Online Fair Allocation of Perishable Resources

Siddhartha Banerjee[1], Chamsi Hssaine[2], and Sean R. Sinclair[1]

[1]School of Operations Research and Information Engineering, Cornell University
[2]Amazon

Abstract

We consider a practically motivated variant of the canonical online fair allocation problem: a decision-maker has a budget of resources to allocate over a fixed number of rounds. Each round sees a random number of arrivals, and the decision-maker must commit to an allocation for these individuals before moving on to the next round. In contrast to prior work, we consider a setting in which resources are *perishable* and individuals' utilities are potentially non-linear (e.g., goods exhibit complementarities). The goal is to construct a sequence of allocations that is *envy-free* and *efficient*. We design an algorithm which takes as input (*i*) a prediction of the perishing order, and (*ii*) a desired bound on envy. Given the remaining budget in each period, the algorithm uses forecasts of future demand and perishing to adaptively choose one of two carefully constructed *guardrail quantities*. We characterize conditions under which our algorithm achieves the optimal envy-efficiency Pareto frontier. We moreover demonstrate its strong numerical performance using data from a partnering food bank. Our results emphasize the importance of high-fidelity predictions of future perishing, all the while highlighting classes of perishable resources for which a decision-maker cannot hope to achieve classical notions of fairness.

Contact author: Sean R. Sinclair (srs429@cornell.edu)

Contents

1 Introduction

Online resource allocation under uncertainty is one of the canonical problems in sequential decision-making/control. While most work in this area focuses on maximizing a single objective (eg. throughput or revenue), more recent work has started looking at fundamental tradeoffs in multi-objective settings. An important example of this is in settings where the principal wants to ensure some form of equitable or *fair* division of the resources while maximizing a primary objective. In particular, while fairness concerns have long been central in the networking and scheduling literature (Kelly et al., 1998; Bonald et al., 2006; Ghodsi et al., 2011; Joe-Wong et al., 2013; Grosof et al., 2019), more recent work has started characterizing the fundamental tradeoffs between fairness and overall utility in different allocation problems (Bertsimas et al., 2011; Donahue and Kleinberg, 2020; Lien et al., 2014; Manshadi et al., 2021; Sinclair et al., 2022; Hassanzadeh et al., 2023). Getting such detailed tradeoffs is useful, as it allows a system designer to understand and choose their desired operating point, balancing the loss in efficiency and equity (or *envy*). All of these works, however, have only considered *linear utility functions*, and *non-restless constraints* (where budgets only change when resources are allocated). This limits their use in many situations of interest, such as the following:

- A food bank often has fixed replenishment intervals (weekly/monthly), and in between two replenishments, has to ration its stocks while serving requests. The food bank manager wants to ensure that allocations are *equitable* across individuals, despite uncertainty in future demand. Unfortunately, as is often the case with food pantries, many food items are perishable, and spoil if not allocated early enough. Moreover, there is often a mix of expiration dates between items, as donations from companies and charities tend to be fresher, while donations from grocery retailers tend to contain near-perishing goods. These expiration dates can serve as a natural schedule according to which items are allocated.
- During the COVID-19 pandemic, the U.S. government tasked itself with distributing ventilators, vaccines, and the antiviral drug Remdesivir (Lupkin, 2020). Each month, states were given a fixed amount of these *complementary* resources. While a primary goal was to distribute as many of these goods as possible (i.e., efficiency), an additional constraint was equitable access on a national level. Achieving these goals proved challenging: between end of 2020 and mid-2022, over 82 million COVID-19 vaccine doses went to waste due to miscalibration of allocations to state's uncertain demands and expiration of multi-vial doses (Eaton, 2022).
- A federated cloud service such as the inter-university Aristotle Cloud Federation (Aristotle Cloud Federation Project, 2022) comprises computer servers contributed by different participating providers (university computing centers). The aim of the cloud scheduler is to adjust the amount of resources allocated to incoming requests while making sure that future requests are guaranteed a minimum level of service (Ghodsi et al., 2011). Moreover, different service providers may remove their unused servers over time to serve internal requests. The scheduler needs to anticipate this *exogenous depletion* while serving its requests.

1.1 Our Contributions

The basic model we study follows the setting considered in Sinclair et al. (2022) (and generalizing single-item allocation models of Lien et al. (2014); Manshadi et al. (2021)): a decision-maker starts with some budget of resources that they must allocate over T discrete rounds to a stream of online arrivals. In each round, arriving agents are drawn from a known distribution (which may be time-varying). Each agent is characterized by a given type, with each type associated with a utility function, which (unlike existing work), we allow to be non-linear in the allocations. More importantly, rather than assuming resource budgets remain fixed unless we allocate them, we consider settings where the resources are *perishable*, and as a result, the budgets are exogenously depleted over time (also according to a known distribution), even if we do not allocate. Our goal is to find a policy that trades off between two metrics:

1. *Envy* – Maximum over agents (in hindsight) of the difference between the utility they receive, and their utility under the allocation received by any other agent; this serves as a measure of (un)fairness.

2. *Inefficiency* – Amount of goods leftover at the end of the time horizon.

We first consider the classical setting of additive utilities. Contrary to prior models, in the perishable resource setting we are faced with two additional challenges. On the one hand, there exist fundamental barriers to what a controller can achieve for general perishing processes. To see this, consider an extreme scenario in which all items perish at the end of the first round. Clearly, there is no hope of achieving low envy in such a setting since future demand cannot be satisfied. In our first main contribution, we identify a class of perishing distributions, *offset-expiring processes* (see Definition 3.2), for which one can meaningfully consider classical notions of envy. For such processes, an important algorithmic challenge remains: an algorithm must have an accurate prediction of *how many* resources will perish in future rounds, but also the *order* in which they perish. This timing aspect of the allocation is the key point of departure from previous models: since items can no longer be allocated after they perish, the algorithm must judiciously choose between the remaining units in each period and allocate those that are likely to perish in subsequent periods to minimize inefficiency.

In our second main contribution, we propose a "guardrail"-style algorithm, Perishing-Guardrail, that takes as input (i) a prediction of the order according to which items perish, (ii) a desired upper bound on envy L_T, and (iii) a high probability value δ. Given these inputs, it computes exactly two allocations (or *guardrails*): high-probability upper and lower bounds on the fair solution in hindsight *without perishing*,that are also within L_T of one another. In each round the algorithm chooses which of the two to guardrails to allocate to each individual, cautiously doing so via the construction of pessimistic forecasts of future arrivals and perishing. Moreover, it allocates according to the input prediction of the perishing order. With this, we have the following main result.

Informal Theorem 1 (see Theorem 4.1)**.** *Suppose the perishing process is offset-expiring. Given a "good enough" prediction of the perishing order,* Perishing-Guardrail *with parameter* L_T *achieves with high probability:*

$$\text{Envy} \lesssim L_T, \qquad \Delta_{\textit{efficiency}} \lesssim \min\{\sqrt{T}, 1/L_T\}.$$

The assumption that the principal has access to a "good" prediction of the perishing order indeed holds in most practical settings — for example, food banks can use the expiry date of canned goods, or the ripeness of fruits and vegetables to determine which are at most risk of spoiling. Moreover, a significant contribution from a modeling perspective is that the class of perishing processes we consider extends far past current state-of-the-art models of online resource allocation of perishable resources, which are limited to *deterministic* or i.i.d. perishing times (Perry, 1999).

Our result then identifies a class of perishable resources for which the envy-efficiency trade-off previously identified in the literature is not fundamentally affected. The intuition behind this is that, when resources perish earlier than expected, they adversely affect both envy and waste; in contrast, uncertainty in the number of arrivals affects the two in opposite directions. Hence, uncertainty in arrivals is the only aspect of this setting that contributes to the trade-off. That said, the design and analysis of an algorithm that is able to account for uncertainty in future perishing constitutes a significant departure from past work. This is due to the fact that though the perishing times are exogenous, the quantity that any good algorithm must forecast is the amount of *unallocated* goods that perish in the future; this quantity is inherently dependent on future decisions made by the algorithm, and is as a result *endogenous*. The key insight that allows us to tackle this challenge is the coupling of our algorithm's sequence of decisions to a tractable "slow" allocation process without perishing; in this process, the number of goods that perish constitute a high-probability upper bound on our algorithm's waste.

Motivated by the limiting nature of the additive preference assumption, we then extend our results to the class of utility functions that are concave, one-positively homogeneous, and non-decreasing. This class of functions subsumes additive utilities, in addition to the popular classes of functions such as Leontief and Cobb-Douglas utilities, which respectively model complementarities and decreasing returns to scale. For such functions, we show that a careful modification of the guardrail construction achieves the envy-efficiency Pareto frontier. The key difficulty in establishing this result is in deriving sharp bounds on the sensitivity of the Eisenberg-Gale (EG) program (Eisenberg, 1961) upon which the guardrail construction relies. In particular, showing that solutions to this program are robust to mild perturbations in the arrival process relies on a property known as *competitive monotonicity*, previously known only for the setting of weak gross substitutes (Jain and Vazirani, 2010).

1.2 Paper Organization

We next survey related work in Section 2. In Section 3, we present the general online resource allocation problem, as well as introducing the perish model and fairness notions of interest. We then design and analyze a guardrail-based algorithm for the setting with perishable goods in Section 4. In Section 5 we extend our results to the setting of non-linear utilities. We conclude by comparing the numerical performance of our algorithms against state-of-the-art benchmarks in Section 6.

2 Related Work

Fairness in resource allocation has a long history in the economics and computation literature, beginning with Varian's seminal work (Varian, 1974, 1976). We highlight the most closely related works below, especially as they relate to *online* fair allocation; see Aleksandrov and Walsh (2019b) for a comprehensive survey.

Fair allocation without perishable resources. We first detail various models and objectives considered in settings without perishable goods. There exists a long line of work in which the *resource* becomes available to the decision-maker online, whereas agents are fixed (Benade et al., 2018; Aleksandrov et al., 2015; Mattei et al., 2018, 2017; Aleksandrov and Walsh, 2019a; Banerjee et al., 2020; Bansal et al., 2020; Bogomolnaia et al., 2021; He et al., 2019; Aziz et al., 2016; Zeng and Psomas, 2020). These models lie in contrast to the one we consider, wherein resources are fixed and *individuals* arrive online. Papers that consider this latter setting include Kalinowski et al. (2013), who consider maximizing *utilitarian* welfare with *indivisible* goods, rather than focusing on *fairness* guarantees with *divisible* goods. Gerding et al. (2019) consider a scheduling setting wherein agents have fixed and known arrival and departure times, as well as demand for the resource. A series of papers also consider the problem of fair division with minimal disruptions relative to previous allocations, as measured by a *fairness ratio*, a competitive ratio analog of counterfactual envy in our setting (Friedman et al., 2017; Cole et al., 2013; Friedman et al., 2015). A number of papers also seek to design algorithms with attractive competitive ratios with respect to the Nash Social Welfare objective (Azar et al., 2010; Banerjee et al., 2020), or the max-min objective (Lien et al., 2014; Manshadi et al., 2021).

The above papers situate themselves within the *adversarial*, or worst-case, tradition. A separate line of work considers fair resource allocation in stochastic settings (Donahue and Kleinberg, 2020; Elzayn et al., 2019; Freund and Hssaine, 2021), as we do. The algorithms developed in these papers, however, are *non-adaptive*: they decide on the entire allocation upfront, *before* observing any of the realized demand. In contrast, we consider a model where the decision-maker makes the allocation decision in each round *after* observing the number of arrivals. Freeman et al. (2017) consider a problem in which agents' utilities are realized from an unknown distribution, and the budget resets in each round. They present algorithms for Nash social welfare maximization and discuss some of their properties. Our work is most closely related to (and indeed, builds upon) Sinclair et al. (2022), who first introduced the envy-freeness and efficiency tradeoff we are interested in. We improve upon their results in showing the applicability of the algorithmic guardrail framework to a broader class of utilities, which subsumes additive utilities that they considered in their original paper. Moreover, we consider a model in which goods also *perish* over time, which none of the aforementioned works consider.

Perishable resources. Though online resource allocation of perishable goods has a long history in the operations research literature (see, e.g., Nahmias (2011) for a comprehensive survey of earlier literature), to the best of our knowledge, the question of *fairly* allocating perishable goods has attracted relatively little attention. Perry (1999) and Hanukov et al. (2020) analyze FIFO-style policies for efficiency maximization in inventory models with Poisson demand and deterministic or Poisson perishing times. Motivated by the problem of electric vehicle charging, Gerding et al. (2019) consider an online scheduling problem where agents arrive and compete for a perishable resource which spoils at the end of every period, and as a result must be allocated at every time step. They consider a range of objectives, including: maximum total resource allocated, maximum number of satisfied agents, as well as envy-freeness. Bateni et al. (2022) similarly consider a setting wherein goods perish immediately. Our paper, in contrast, considers *stochastic* perishing over the course of multiple rounds. Alijani et al. (2020) similarly consider a setting with stochastic perishing; in the problem they consider, a decision-maker seeks to sell perishable items to a stream of buyers in order to maximize social welfare. They show that, when items have independent perishing times satisfying the monotone hazard rate condition, the competitive ratio of any policy is lower bounded by a constant greater than one. This negative result is in line with our discussion regarding i.i.d. perishing times representing the worst-case for the decision-maker. Also related is recent work on stochastic matching with unknown arrivals and abandonments. For instance, Aouad and Saritaç (2020) design constant-factor approximations for a setting where agents of different types arrive according to a Poisson process and abandon the system once their exponentially distributed sojourn time elapses, and the decision-maker seeks to maximize cumulative rewards.

3 Preliminaries

In this section, we state the most general form of the model, relegating specifics for the applications we consider to subsequent sections.

Notation. We use $\mathbb{R}_+$ to denote the set of non-negative reals, $\|X\|_\infty = \max_{i,j} |X_{i,j}|$ to denote the matrix maximum norm, and cX to denote entry-wise multiplication for a constant c. When comparing vectors, we use $X \leq Y$ to denote that each component $X_i \leq Y_i$. Finally, given $n \in \mathbb{N}^+$, we let $[n] = \{1, \dots, n\}$.

3.1 Model

We consider a decision-maker who, over T distinct rounds, must divide K divisible resources among a population of individuals. The decision-maker has an initial fixed budget of $B_k \in \mathbb{R}$ divisible units (also referred to as *items*, or *goods*) of each resource $k \in [K]$. Let $B = (B_k)_{k \in [K]}$. We often abuse notation and use $[B] = \{1, \dots, \sum_k B_k\}$ to denote the set of all items.

Demand model. In each round $t \in [T]$, a random number of individuals arrives, each requesting a share of the resources. Each individual is characterized by their type $\theta \in \Theta$, with $|\Theta| < \infty$. Specifically, each type $\theta \in \Theta$ is associated with a *known* utility function $u(x, \theta) : \mathbb{R}^K \times \Theta \mapsto \mathbb{R}_+$, for a given allocation $x \in \mathbb{R}^K$ of resources.

We let $N_{t,\theta}$ denote the number of type θ arrivals in round t, with $N_{t,\theta}$ drawn independently from a known distribution $\mathcal{F}_{t,\theta}$. For a fixed vector of arrivals $(N_{t,\theta})_{t \in [T], \theta \in \Theta}$, we will often use $N_{\geq t,\theta}$ to denote the total number of individuals who arrived in round t and afterwards (i.e., $N_{\geq t,\theta} = \sum_{t'=t}^{T} N_{t',\theta}$). We similarly let $N_{\leq t,\theta} = \sum_{t'=1}^{t} N_{t',\theta}$, and use $N_{[t,t'),\theta}$ to denote arrivals of type θ between rounds t and t'. We use N_θ to denote the *total* number of type θ individuals across all rounds (i.e., $N_\theta = \sum_{t \in [T]} N_{t,\theta}$), and let $N = \sum_{t \in [T]} \sum_{\theta \in \Theta} N_{t,\theta}$. For ease of notation, we use similar subscripting for random variables throughout this paper. Whenever the θ subscript is omitted, it is assumed that we are considering the aggregate quantity, summed over all θ.

Perishing model. Each unit of resource $b \in [B]$ is associated with a perishing time $T_b \in \mathbb{N}^+$ drawn from a *known* distribution. We assume items' perishing times are independent of one another, as well as of the arrival process, and that perishing occurs at the *end* of each round, after items have been allocated to individuals. For $t \in [T]$, we let $P_t = \sum_{b \in [B]} \mathbb{1}\{T_b = t\}$ denote the number of units of resource perishing in period t, and define $P_{<t} = \sum_{\tau < t} P_t$ to denote the number of resources that perished before period t; we use ν_t to denote its expected value.

Objective. The goal is to design a *fair* online algorithm that observes the number of arrivals of each type $N_{t,\theta}$ *at the beginning* of round t and determines (i) the *amount* of resource k to give out to each type $\theta \in \Theta$, for all $k \in [K]$, and (ii) the *order* in which to allocate these resources. Let $X^{alg} \in \mathbb{R}^{T \times |\Theta| \times K}$ be the sequence of allocations determined by the algorithm. We assume that, for $t \in [T], \theta \in \Theta$, any algorithm allocates $X_{t,\theta}$ uniformly across all $N_{t,\theta}$ individuals.

3.1.1 Additional assumptions and notation.

For technical simplicity, we begin by considering the classical setting with additive utilities: $u(X_{t,\theta}, \theta) = \sum_k w_{\theta k} X_{t,\theta,k}$ for some collection of weights $\{w_{\theta k}\}_{k \in [K]}$, and $K = 1$ (thus, $w_\theta = 1$ without loss of generality); we show how to extend our results to a more general class of utility functions and $K \geq 1$ in Section 5.

For all $t \in [T], \theta \in \Theta$, we assume $N_{t,\theta} \geq 1$ almost surely. This is for ease of exposition; our results continue to hold (up to constants) as long as $\mathbb{P}(N_{t,\theta} = 0)$ does not scale with T. We also assume that $\mathbb{E}[N] = \Theta(T)$. We define $\beta_{avg} = \frac{\sum_k B_k}{\mathbb{E}[N]}$ to be the average resource per individual, and assume $\beta_{avg} \in \Theta(1)$. We moreover let $\sigma_{t,\theta}^2 = \text{Var}(N_{t,\theta}) < \infty$, and assume $\rho_{t,\theta} = |N_{t,\theta} - \mathbb{E}[N_{t,\theta}]| < \infty$ almost surely. Finally, we let $\mu_{\max} = \max_{t,\theta} \mathbb{E}[N_{t,\theta}], \sigma_{\min}^2 = \min_{t,\theta} \sigma_{t,\theta}^2, \sigma_{\max}^2 = \max_{t,\theta} \sigma_{t,\theta}^2$, and $\rho_{\max} = \max_{t,\theta} \rho_{t,\theta}$. We summarize all notation in Appendix A.

3.2 Notions of Fairness and Efficiency

Before defining our notions of fairness and efficiency for the *online* setting, in which the decision-maker observes the arrivals $(N_{t,\theta})_{\theta \in \Theta}$ *at the beginning* of each round $t \in [T]$, we introduce standard notions of fairness and efficiency in the *offline* setting, in which the decision-maker knows the entire vector of arrivals $(N_{t,\theta})_{t \in [T], \theta \in \Theta}$ at the beginning of time.

3.2.1 Fairness and Efficiency in Offline Allocations

The notion of *offline* fairness we consider is that of *Varian Fairness* (Varian, 1974), widely used in the operations research and economics literature.

Definition 3.1 (Fair Allocation)**.** *Given types Θ, number of individuals of each type $(N_{t,\theta})_{t \in [T], \theta \in \Theta}$, perishing times $(T_b)_{b \in [B]}$, and utility functions $(u(\cdot, \theta))_{\theta \in \Theta}$, an allocation $X = \{X_{t,\theta} \in \mathbb{R}_+^K \mid \sum_{t=1}^T \sum_{\theta \in \Theta} N_{t,\theta} X_{t,\theta} \leq B\}$ is* fair *if it simultaneously satisfies the following:*

1. *Envy-Freeness (*EF*): $u(X_{t,\theta}, \theta) \geq u(X_{t',\theta'}, \theta)$ for all $t, t' \in [T], \theta, \theta' \in \Theta$*
2. *Pareto Efficiency (*PE*): Consider feasible allocations $Y \in \mathbb{R}^{T \times |\Theta| \times K}$, $X \in \mathbb{R}^{T \times |\Theta| \times K}$ such that $Y \neq X$, and $u(Y_{t,\theta}, \theta) > u(X_{t,\theta}, \theta)$ for $t \in [T], \theta \in \Theta$. Then, there exists $t' \in [T], \theta' \in \Theta$ such that $u(Y_{t',\theta'}, \theta') < u(X_{t',\theta'}, \theta')$.*
3. *Proportionality (*PROP*): For all $t \in [T], \theta \in \Theta$, $u(X_{t,\theta}, \theta) \geq u(B/N, \theta)$.*

Contrary to the setting without perishable resources, a fair solution need not exist for arbitrary perishing processes, as illustrated in Section 1. For example, suppose all resources expire at the end of the first day; demand for days $t \geq 2$ cannot be satisfied no matter the allocation decisions. In such settings, re-defining envy to be forward-looking would be more appropriate, since customers have no reason to "envy" products that perished before they arrive. We leave the treatment of such important modeling questions to future work, and instead restrict our attention to consider "offset-expiring" processes which are such that B/N, the hindsight optimal allocation *without* perishing (Sinclair et al., 2022), remains optimal.

Definition 3.2. *A perishing process $(P_t)_{t \in [T]}$ is **offset-expiring** if $P_{<t} \leq \frac{B}{N} N_{<t}$, for all $t \in [T]$.*

Such processes model settings wherein the number of perished resources up to any round t is smaller than what the optimal hindsight allocation decision would have allocated up to that time. In Section 4 we give a sufficient condition on the joint distribution over perishing and arrivals for the process to be offset-expiring with high probability. One would expect this condition to hold in practice; as noted in Section 1, food banks typically have an idea of goods' expiry dates, and may even refuse to accept items that have a high risk of spoiling before the next replenishment.

We next present a convex program, termed the offline *perishing Eisenberg-Gale (EG) program*, which we will show produces fair solutions for offset-expiring perishing processes.

$$\begin{aligned} \max_{X \in \mathbb{R}^{T\times|\Theta|}} \quad & \sum_{t\in[T]}\sum_{\theta\in\Theta} N_{t,\theta} \log u(X_{t,\theta},\theta) \\ \text{s.t.} \quad & \sum_{\tau\geq t}\sum_{\theta} N_{\tau,\theta}X_{\tau,\theta} \leq B - \sum_{\tau<t} P_\tau \quad \forall t\in[T] \\ & X_{t,\theta}\geq 0 \quad \forall t\in[T], \theta\in\Theta \end{aligned} \tag{EG-P}$$

(EG-P) is a variant of the well-known EG program (Eisenberg, 1961), for which solutions are well-known to satisfy the three fairness properties under a special class of utility functions (including additive utilities). In particular, (EG-P) accounts for perishable resources through its first set of constraints which enforces that, for each period, the amount allocated to individuals in subsequent periods cannot exceed the number of items that haven't perished yet. Note that (EG-P) is an *aggregate* optimization problem which does not encode the order in which items must be allocated. However, it is easy to show that the optimal value of this relaxation is exactly equal to that of a more granular program wherein timing is indeed encoded, under the optimal hindsight timing rule which allocates items according to the known perishing time; we thus omit the proof of this fact. Moreover, when the utility functions are concave, solutions to (EG-P) are efficiently computable.

The following proposition implies that, under offset-expiring processes, (EG-P) produces a fair solution. We defer its proof to Appendix B.

Proposition 3.3. *Suppose $(P_t)_{t\in[T]}$ is offset-expiring. Then, $X^{opt}_{t,\theta} = B/N$ for all $t\in[T], \theta\in\Theta$ solves (EG-P) and is fair.*

In contrast to the *online* problem we are interested in, which only has access to $(N_{t,\theta})_{\theta\in\Theta}$ at the beginning of round t, the perishing EG program is *offline*, i.e., it has access to *all* arrivals at the beginning of time. Unfortunately, though computing fair allocations is possible in the offline setting for this class of perishing processes, Sinclair et al. (2022) showed that no *online* algorithm can simultaneously achieve the three desired properties of envy-freeness, Pareto efficiency and proportionality, even in a simple setting without perishable resources. This then motivates the goal of finding *approximately* fair allocations in the online setting.

3.2.2 Approximate Fairness and Efficiency in Online Allocations

We formally define the notions of *online* fairness we are interested in, first introduced in Sinclair et al. (2022). Let $X^{alg}\in\mathbb{R}^{T\times|\Theta|\times K}$ be the allocations determined by an arbitrary online algorithm.

Definition 3.4 (Counterfactual Envy, Hindsight Envy, Efficiency, and Proportionality)**.** *Given individuals with types Θ, sizes $(N_{t,\theta})_{t\in[T],\theta\in\Theta}$, resource budgets $(B_k)_{k\in[K]}$, and perishing times $(T_b)_{b\in[B]}$, for any online allocation $(X^{alg}_{t,\theta})_{t\in[T],\theta\in\Theta}\in\mathbb{R}^K$, we define:*

- Counterfactual Envy*: The counterfactual distance of X^{alg} to envy-freeness as*

$$\Delta_{EF} \triangleq \max_{t\in[T],\theta\in\Theta} \|u(X^{alg}_{t,\theta},\theta) - u(X^{opt}_{t,\theta},\theta)\|_\infty \tag{1}$$

  where X^{opt} is a solution to the offline EG program (EG-P).

- Hindsight Envy*: The hindsight distance of X^{alg} to envy-freeness as*

$$\textsc{Envy} \triangleq \max_{t,t'\in[T]^2,\theta,\theta'\in\Theta^2} u(X^{alg}_{t',\theta'},\theta) - u(X^{alg}_{t,\theta},\theta). \tag{2}$$

- Efficiency*: The distance to efficiency as*

$$\Delta_{\textit{efficiency}} \triangleq \sum_{k \in K} \left(B_k - \sum_{t \in [T]} \sum_{\theta \in \Theta} N_{t,\theta} X^{alg}_{t,\theta,k} - \sum_{t \in [T]} \mathrm{PUA}_{t,k} \right), \tag{3}$$

where $\mathrm{PUA}_{t,k} = \sum_{b \in B_k} \mathbb{1}\{T_b = t, b \notin \mathcal{A}_t\}$ *denotes the amount of resource k that perished at the end of round t before having been allocated.*

- Hindsight Proportionality*: The hindsight distance of X^{alg} to proportionality as*

$$\Delta_{\textit{prop}} \triangleq \max_{t \in [T], \theta} u(B/N, \theta) - u(X^{alg}_{t,\theta}, \theta). \tag{4}$$

At a high level, counterfactual envy Δ_{EF} can be viewed as the online algorithm's distance to fairness *in hindsight*. Though this latter envy-freeness metric is with respect to an offline solution, hindsight envy measures how differently the online algorithm treats any two individuals, across types and time. Finally, the efficiency of the online algorithm $\Delta_{\textit{efficiency}}$ measures how *wasteful* the algorithm was in hindsight. Namely, if at the end of the horizon, the decision-maker has a large number of unallocated goods remaining, she could have gone back and re-distributed these items, thereby increasing their utilities. The reader may notice that this definition of distance to efficiency does not exactly mirror the offline notion of Pareto efficiency defined in the previous section. However, it is easy to see that under strict monotonicity of utility functions, a distance to efficiency of zero is a necessary precondition to Pareto efficiency.

Note that this definition of waste does not include waste incurred from perishing; this is due to the fact that, even with *perfect knowledge* of the number of arrivals and the order in which items perish, one can construct offset-expiring processes such that a non-trivial amount of spoilage occurs under any envy-free allocation (see Section 4.3 for further discussion of this). Defining inefficiency in this way avoids penalizing our algorithm for unavoidable spoilage; we leave the investigation of tight bounds on *total* waste to future work.

These metrics are, in a sense, at odds with each other, which marks the inherent difficulty of this problem. To see this, consider the following two extreme scenarios. On the one hand, an algorithm can trivially achieve a hindsight envy of zero by allocating nothing to individuals in any round; this, however, would result in both high *counterfactual* envy and proportionality, in addition to maximal inefficiency. On the other hand, a distance to efficiency of zero can trivially be satisfied by exhausting the budget in the initial round, at a cost of maximal envy as individuals arriving at later rounds would be envious of the allocation given in the first. Sinclair et al. (2022) formalized this tension for the additive utility setting without perishable resources via the following lower bounds, which we re-state here for completeness.

Theorem 3.5 (Theorems 1 and 2, Sinclair et al. (2022))**.** *Under any arrival distribution satisfying the assumptions outlined above, there exists a problem instance with $K = 1$, additive utilities, and no perishing, such that any algorithm must incur $\Delta_{EF} \gtrsim \frac{1}{\sqrt{T}}$. Moreover, any algorithm that achieves $\Delta_{EF} \leq L_T = o(1)$ or $\textsc{Envy} \leq L_T = o(1)$ must also incur waste $\Delta_{\textit{efficiency}} \gtrsim \min\{\sqrt{T}, 1/L_T\}$.*

This lower bound also holds for the more general setting we consider, as even with arbitrary utilities the optimal fair allocation in the case of a single resource is $X^{opt}_{t,\theta} = \frac{B}{N}$, for $t \in [T], \theta \in \Theta$, where $N = \sum_{t,\theta} N_{t,\theta}$. This follows from the fact that solutions to the EG program are envy-free, and the *unique* envy-free allocation and Pareto-efficient solution is B/N in the case of a single resource.

Since settings without perishable resources are a special case of our setting (e.g., a perishing process with $T_b > T$ a.s., for all $b \in [B]$), this lower bound holds in our case; the goal will then be to design algorithms that achieve this lower bound with high-probability.

Attempts to naively utilize algorithms achieving this lower bound for the setting without perishable resources face an obvious first obstacle in our setting. Any algorithm needs to properly use the perishing prediction as well as leverage distributional information on the perishing process to predict future perishing resources. This is difficult due to the interdependence between future perishing and future allocations; not only does the algorithm's decision today impact future perishing (i.e., allocating many, resp. few, goods today results in potentially less, resp. more, perishing tomorrow), but the algorithm's decision today must consider such effects of all future decisions to construct a prediction of future perishing. We present an algorithm that tackles these subtleties and achieves the optimal envy-efficiency tradeoff in the following section.

4 Online Fair Allocation with Perishing Resources

4.1 The Algorithm

Our algorithm, Perishing-Guardrail, takes as input (i) a desired bound on envy L_T, (ii) a pre-specified *allocation schedule* σ (also referred to as an *ordering* or *priority list*), according to which it allocates items, and (iii) a high probability parameter δ. The allocation schedule functionally acts as a prediction of the order in which items will perish; though our theoretical results hold for a variety of schedules, practically speaking "good" schedules allocate the items that are most likely to perish early on, and defer items that are unlikely to perish to later rounds.

Using a high-probability upper bound on the number of arrivals throughout the time horizon, denoted by $\overline{N}$, our algorithm computes two allocations, $\underline{X}$ and $\overline{X}$, that it will restrict itself to use as long as there is sufficient budget remaining:

$$\begin{cases} \underline{X} & = \frac{B}{\overline{N}}\left(1 - \frac{c}{\sqrt{T}}\right), \quad c = (1 + \sqrt{3\log(2T/\delta)})/\beta_{avg} \\ \overline{X} & = \underline{X} + L_T. \end{cases}$$

$\underline{X}$, the *lower guardrail*, is then a high-probability lower bound on the hindsight fair solution $X_t^{opt} = B/N$, and $\overline{X}$, the *upper guardrail*, is constructed to be L_T away.

In order to decide when the algorithm allocates the lower versus upper guardrail in periods where it has enough budget remaining, it checks if it can allocate the higher allocation $\overline{X}$ assuming (i) a high-probability upper bound on the number of future arrivals guaranteeing them the lower guardrail allocation, and (ii) a high-probability upper bound on future perishing. If so, it chooses this higher allocation $\overline{X}$, otherwise, it allocates $\underline{X}$. Once the allocation is decided, the algorithm allocates the resources in order according to the allocation schedule σ.

Under this algorithm, envy is managed by the construction of $\underline{X}$ and $\overline{X}$ being L_T away, and is only ever additionally incurred if the algorithm runs out of budget, in which case agents later on in the time horizon will receive nothing. On the other hand, waste is incurred if the algorithm allocates too little relative to the budget.

Before formally presenting our algorithm, we introduce notation that will be of use throughout.

Notation. For ease of notation, for $t < t'$ we let $\underline{N}_{[t,t']} = \mathbb{E}[N_{[t,t']}] - \text{Conf}^N_{t,t'}$, $\overline{N}_{[t,t']} = \mathbb{E}[N_{[t,t']}] + \text{Conf}^N_{t,t'}$ be high-probability upper and lower bounds on realized demand between t and t' for appropriately defined confidence terms. For $t \in [T]$, let $\mathcal{A}_t$ denote the set of items allocated in round t, and PUA_t the quantity of unallocated items that perished at the end of round t. For $i \in [B]$, $\sigma(i)$ denotes the ith-ranked item in the allocation schedule. For any two items b, b' such that b comes (weakly) before b' in σ, we write $b \preceq_\sigma b'$. Then, for any subset $S \subseteq [B]$ of items, $\textsc{Prec}(b \mid S) = \sum_{b' \in S} \mathbb{1}\{b' \preceq_\sigma b\}$ represents the number of items in S that are ranked (weakly) before b according to σ. It will often be useful for us to consider a "slow" resource consumption sample path, in which $\underline{N}_t \underline{X}$ items are allocated in each period $t \in [T]$, and no items perish ahead of time. Fixing $t \in [T]$, $S \subseteq [B]$ and $b \in S$, we let $\tau_b(t \mid S)$ be the period in which b *would* have been entirely allocated after t under a slow resource consumption sample path. Formally, $\tau_b(t \mid S) = \inf\{t' \geq t : \underline{N}_{[t:t']} \underline{X} > \textsc{Prec}(b \mid S)\}$. Under this slow process, we moreover define the remaining set of items in period t to be $[\overline{B}_t] = \{\sigma(\lfloor \underline{N}_{<t} \underline{X} \rfloor), \ldots, \sigma(B)\}$. With this notation in hand we finally define $\eta_t = \sum_{b \in [\overline{B}_t]} \mathbb{P}\big(t \leq T_b < \min\{T, \tau_b(t \mid [\overline{B}_t])\}\big)$, the expected number of perished resources from t onwards under this slow process. See Table 1 for a summary of this notation.

We now present our algorithm, $\textsc{Perishing-Guardrail}$, in Algorithm 1.

ALGORITHM 1: $\textsc{Perishing-Guardrail}$

Input: Budget $B = B_1^{alg}$, allocation schedule σ, envy parameter L_T, confidence terms with respect to the arrival process $(\textsc{Conf}^N_{t,t'})_{t,t' \in [T], \theta \in \Theta}$ and perishing inputs $(\eta_t)_{t \in [T]}$
Output: An allocation $X^{alg} \in \mathbb{R}^{T \times |\Theta|}$
Compute $\underline{X} = \frac{B}{\mathbb{E}[N] + \textsc{Conf}^N_{0,T}} \left(1 - \frac{c}{\sqrt{T}}\right)$ for $c = (1 + \sqrt{3 \log(2T/\delta)})/\beta_{avg}$ $\quad \overline{X} = \underline{X} + L_T$
for $t = 1, \ldots, T$ **do**
  Compute $\overline{P}_t = \eta_t + \sqrt{3 \log(2T/\delta) \eta_t}$ `// Compute ''worst-case'' perishing`
  if $B_t^{alg} < N_t \underline{X}$ **then** `// insufficient budget to allocate lower guardrail`
    Set $X_{t,\theta}^{alg} = \frac{B_t^{alg}}{N_t}$ for each $\theta \in \Theta$. Allocate items $b \in [B_t^{alg}]$ according to σ.
  else if $B_t^{alg} - N_t \overline{X} \geq \underline{X}(\mathbb{E}[N_{>t}] + \textsc{Conf}^N_{t,T}) + \overline{P}_t$ **then** `// use upper guardrail`
    Set $X_{t,\theta}^{alg} = \overline{X}$ for each $\theta \in \Theta$. Allocate items $b \in [B_t^{alg}]$ according to σ.
  else `// use lower guardrail`
    Set $X_{t,\theta}^{alg} = \underline{X}$ for each $\theta \in \Theta$. Allocate items $b \in [B_t^{alg}]$ according to σ.
  Update $B_{t+1}^{alg} = B_t^{alg} - N_t X_t^{alg} - \text{PUA}_t$
end
return X^{alg}

4.2 Performance Guarantee

In our main result, we show that under offset-expiring models, if (i) perishing is "slow," and (ii) the prediction of the perishing order σ is high-fidelity (in a sense made formal below), $\textsc{Perishing-Guardrail}$ achieves the optimal envy-efficiency tradeoff with high probability. We defer a discussion of the conditions under which our algorithm enjoys its guarantees to Section 4.3.

Theorem 4.1. *For $t' > t$, let $\textsc{Conf}^N_{t,t'} = \sqrt{2(t'-t)|\Theta| \rho_{max}^2 \log(2T^2/\delta)}$, and $\textsc{Conf}^P_t = \sqrt{3\nu_t \log(2T/\delta)}$, and define*

$$\overline{P}_t = \eta_t + \sqrt{3 \log(2T/\delta) \eta_t}.$$

Suppose the following conditions hold, for all $t \in [T]$:

(1) $\mathbb{E}[P_{<t}] + \textsc{Conf}^P_t \leq \frac{\mathbb{E}[N_{<t}] - \textsc{Conf}^N_{0,t-1}}{\mathbb{E}[N] + \textsc{Conf}^N_{0,T}} B$ *and*

(2) $\eta_t \leq \sqrt{T - t}$.

Then with probability at least $1 - 4\delta$*,* Perishing-Guardrail *achieves:*

$$\Delta_{EF} \lesssim \max\{1/\sqrt{T}, L_T\} \qquad \Delta_{\textit{efficiency}} \lesssim \min\{\sqrt{T}, 1/L_T\} \qquad \textsc{Envy} \lesssim L_T \qquad \Delta_{prop} \lesssim \max\{1/\sqrt{T}, L_T\},$$

where $\lesssim$ *drops poly-logarithmic factors of* T*,* $\log(1/\delta)$*,* $o(1)$ *terms, and absolute constants.*

Before proving our main technical result, we discuss salient aspects of the theorem statement. Condition (1) of Theorem 4.1 is a sufficient condition on the joint distribution over perishing and arrivals for the process to be offset-expiring with high probability, constructed via straightforward concentration arguments. This condition isn't expected to be restrictive in practice, as discussed in Section 3. We also note that the dependence on δ can be removed by taking $\delta = O(1/T^2)$; in this case, the derived bounds hold in expectation.

We next state the main building blocks upon which the theorem relies, deferring their proofs to Appendix C.

The theorem relies on a careful analysis of both the allocations decided by Perishing-Guardrail and the quantity of unallocated items that perished at the end of each round, under event $\mathcal{E}$, defined to be the intersection of the following three events:

1. $\mathcal{E}_N = \{|N_{(t,t']} - \mathbb{E}\big[N_{(t,t']}\big]| \leq \textsc{Conf}^N_{t,t'} \ \forall\ t, t' > t\ \}$
2. $\mathcal{E}_P = \{|P_{<t} - \mathbb{E}[P_{<t}]| \leq \textsc{Conf}^P_t \ \forall\ t \in [T]\}$
3. $\mathcal{E}_{\overline{P}} = \{\overline{P}_t \geq \mathrm{PUA}_{\geq t} \forall\ t \in [T]\}$, where $\mathrm{PUA}_{\geq t}$ represents the quantity of unallocated items that perished between the end of round t and the end of round $T-1$.[1]

In words, $\mathcal{E}$ represents the event that the arrival and perishing processes fall close to their respective means, and moreover that $\overline{P}_t$ is indeed a pessimistic estimate of the unallocated goods that perish in the future. The following lemma implies that it suffices to restrict attention to $\mathcal{E}$.

Lemma 4.2. *Let* $\mathcal{E} = \mathcal{E}_N \cap \mathcal{E}_P \cap \mathcal{E}_{\overline{P}}$*. Then,* $\mathbb{P}(\mathcal{E}) \geq 1 - 4\delta$*.*

Though the concentration bounds associated with events $\mathcal{E}_N$ and $\mathcal{E}_P$ follow from standard applications of Hoeffding's inequality, the high-probability bound associated with $\mathcal{E}_{\overline{P}}$ presents additional challenges. To upper bound the amount of unallocated resources that perished between t and $T-1$, we must now account for both the uncertainty in arrivals N_t, in addition to the realized order in which resources perished, and relate these two sources of uncertainty to the time at which the algorithm intended to allocate these resources. Establishing this fact hinges upon the careful construction of the "slow" process described above, which decouples future perishing from future allocations to compute $\overline{P}_t$. Note moreover that $\mathcal{E}$ and Condition (2) together imply that the total spoilage under our algorithm is $O(\sqrt{T})$.

Given $\mathcal{E}$, Lemma 4.3 establishes that in every round, the algorithm has enough budget to allocate at least the lower guardrail $\underline{X}$ to all future arrivals.

Lemma 4.3. *Under event* $\mathcal{E}$*,* $B^{alg}_t \geq N_{\geq t}\underline{X}$ *for all* $t \in [T]$*.*

As a result, either $\overline{X}$ or $\underline{X}$ is allocated in every round, thus guaranteeing envy of at most L_T. With that we turn our attention to obtaining a bound on $\Delta_{\textit{efficiency}}$. Note that whenever the algorithm allocates $\underline{X}$, by the threshold condition it must be that the remaining budget is "tight within the confidence radius" of the future demand (hence, low waste). However, since the

[1] Perishing at the end of round T does not "matter" to the performance of our algorithm, in a sense, since anything that is unallocated at the end of round T (perished or not) counts as waste.

confidence radius scales with respect to the number of remaining rounds in the horizon, this is an easier condition to satisfy early on versus later in the horizon. We first address this with the following lemma, which states that there exists a timestep t_0 close to the end of the time horizon for which the algorithm allocates $\underline{X}$ (and allocates $\overline{X}$ for every round afterwards).

Lemma 4.4. *Let t_0 be the last time that $X_t^{alg} \neq \overline{X}$ (or else 0 if the algorithm always allocates according to $\overline{X}$). Then, under event $\mathcal{E}$, for some $\tilde{c} = \widetilde{\Theta}(1)$, $t_0 \geq \max\{0, T - \tilde{c}L_T^{-2}\}$.*

With these lemmas in hand, we prove our main result.

Proof. By Lemma 4.3, the algorithm never runs out of budget under event $\mathcal{E}$, which occurs with probability at least $1 - 4\delta$. In the remainder of the proof, we assume event $\mathcal{E}$ holds.

The following lemma implies that $\underline{X}$ is indeed a good approximation of X^{opt}, given $\mathcal{E}$.

Lemma 4.5. *Given $\mathcal{E}$, $X_{t,\theta}^{opt} = \frac{B}{N}$ for all $t \in [T]$, $\theta \in \Theta$.*

We now proceed with the bounds on envy and efficiency.

Counterfactual Envy: Under $\mathcal{E}$, $\underline{X} \leq \frac{B}{N} = X_{t,\theta}^{opt}$ for all t, θ. We break the proof out into two different cases, depending on $\overline{X}$.

Case 1: $\underline{X} \leq \overline{X} \leq \frac{B}{N}$. By definition:

$$\Delta_{EF} = \frac{B}{N} - \underline{X} \leq \frac{B}{\mathbb{E}[N] - \textsc{Conf}_{0,T}^N} - \frac{B}{\mathbb{E}[N] + \textsc{Conf}_{0,T}^N}\left(1 - \frac{c}{\sqrt{T}}\right)$$
$$= \frac{B}{\mathbb{E}[N]}\left(\frac{1}{1 - \frac{\textsc{Conf}_{0,T}^N}{\mathbb{E}[N]}} - \frac{1 - c/\sqrt{T}}{1 + \frac{\textsc{Conf}_{0,T}^N}{\mathbb{E}[N]}}\right) = \beta_{avg}\left(\frac{1}{1 - \frac{\textsc{Conf}_{0,T}^N}{\mathbb{E}[N]}} - \frac{1 - c/\sqrt{T}}{1 + \frac{\textsc{Conf}_{0,T}^N}{\mathbb{E}[N]}}\right).$$

Using the fact that $\textsc{Conf}_{0,T}^N = \sqrt{2T|\Theta|\rho_{\max}^2 \log(2T^2/\delta)}$ and $\mathbb{E}[N] \in \Theta(T)$, there exists $c_1, c_2 \in \widetilde{\Theta}(1)$ such that, for large enough T, $\left(1 - \frac{\textsc{Conf}_{0,T}^N}{\mathbb{E}[N]}\right)^{-1} \leq \left(1 - c_1/\sqrt{T}\right)^{-1} \leq 1 + 2c_1/\sqrt{T}$ and $\left(1 + \frac{\textsc{Conf}_{0,T}^N}{\mathbb{E}[N]}\right)^{-1} \geq \left(1 + c_2/\sqrt{T}\right)^{-1} \geq 1 - c_2/\sqrt{T}$. Plugging this into the above we have that:

$$\Delta_{EF} \leq \beta_{avg}\left(1 + 2c_1/\sqrt{T} - (1 - c/\sqrt{T})(1 - c_2/\sqrt{T})\right)$$
$$\leq \beta_{avg}(2c_1 + c_2 + c)/\sqrt{T} \lesssim 1/\sqrt{T}.$$

Case 2: $\underline{X} \leq \frac{B}{N} \leq \overline{X}$. We have that:

$$\Delta_{EF} = \max\left\{\frac{B}{N} - \underline{X}, \frac{B}{N} - \overline{X}\right\} \leq \overline{X} - \underline{X} = L_T.$$

Combining these two cases gives that $\Delta_{EF} \lesssim \max\{1/\sqrt{T}, L_T\}$ as needed.

Hindsight Envy: $\textsc{Envy}$ is trivially bounded above by L_T since we have that for any t, t':

$$X_{t'} - X_t \leq |\overline{X} - \underline{X}| = L_T.$$

Efficiency: We next consider the bound on distance to efficiency $\Delta_{\mathit{efficiency}}$. Here we leverage Lemma 4.4, and let t_0 be the last time that $X_t \neq \overline{X}$. This implies that:

$$\begin{aligned}
\Delta_{\mathit{efficiency}} &= B - \sum_t N_t X_t^{alg} - \mathrm{PUA}_{\leq T} \\
&= B_{t_0} + \sum_{t<t_0} N_t X_t^{alg} + \mathrm{PUA}_{<t_0} - \sum_t N_t X_t^{alg} - \mathrm{PUA}_{\leq T} \\
&\leq B_{t_0} - \sum_{t \geq t_0} N_t X_t^{alg} \\
&< N_{t_0}\overline{X} + \overline{N}_{>t_0}\underline{X} + \overline{P}_{t_0} - N_{t_0}\underline{X} - N_{>t_0}\overline{X} \\
&= \underline{X}(\overline{N}_{>t_0} - N_{>t_0}) - (\overline{X} - \underline{X})(N_{>t_0} - N_{t_0}) + \overline{P}_{t_0},
\end{aligned}$$

where the first inequality follows from the fact that $\mathrm{PUA}_{<t_0} \leq \mathrm{PUA}_{\leq T}$, and the second from $X_{t_0}^{alg} = \underline{X}$. Noting that $\underline{X} \leq \beta_{avg}$ and $\overline{N}_{>t_0} - N_{>t_0} \leq 2\mathrm{C{\small ONF}}_{t_0,T}^N$, we have the following upper bound on the first term:

$$\underline{X}(\overline{N}_{>t_0} - N_{>t_0}) \leq \beta_{avg} \cdot 2\sqrt{2(T-t_0)|\Theta|\rho_{max}^2 \log(2T^2/\delta)}.$$

We loosely upper bound the second term by $(\overline{X} - \underline{X})N_{t_0} \lesssim L_T|\Theta|(\mu_{max} + \rho_{max})$. Finally, by Lemma 4.4 we have that $t_0 \geq \max\{0, T - \tilde{c}L_T^{-2}\}$ for some $\tilde{c} \in \widetilde{\Theta}(1)$. Hence, we obtain:

$$\begin{aligned}
\Delta_{\mathit{efficiency}} \leq\ & 2\beta_{avg}\sqrt{2\tilde{c}|\Theta|\rho_{max}^2 \log(2T^2/\delta)} \min\{\sqrt{T}, 1/L_T\} + L_T|\Theta|(\mu_{max} + \rho_{max}) \\
& + (1 + \sqrt{3\log(2T/\delta)})\sqrt{T - t_0} \\
\leq\ & 2\beta_{avg}\sqrt{2\tilde{c}|\Theta|\rho_{max}^2 \log(2T^2/\delta)} \min\{\sqrt{T}, 1/L_T\} + L_T|\Theta|(\mu_{max} + \rho_{max}) + \\
& (1 + \sqrt{3\log(2T/\delta)}) \min\{\sqrt{T}, 1/L_T\},
\end{aligned}$$

where we combined the definition of $\overline{P}_{t_0}$ with the upper bound on $T - t_0$ to finalize the bound.

Proportionality: In this setting, $\Delta_{prop} = \Delta_{EF}$ since $X^{opt} = \frac{B}{N}$, so $\Delta_{prop} = \Delta_{EF} \lesssim \max\{1/\sqrt{T}, L_T\}$. □

4.3 Discussion of Assumptions

The motivation behind the conditions in Theorem 4.1 is similar to the one outlined in Section 3 regarding offset-expiring processes. In practical settings such as food pantries, it must be that a vanishingly small number of goods perish in between replenishments (otherwise, the food bank is being run poorly). For Condition (2), note the natural dependence on t in the upper bound. As t increases, the "worst-case" allocation time $\tau_b(t \mid [\overline{B}_t])$ can only (weakly) decrease. Thus, the likelihood that a good's perishing time falls in $[t, \min\{T, \tau_b(t \mid [\overline{B}_t])\})$ can only decrease as t increases. Putting this together with the fact that the size of the set $[\overline{B}_t]$ itself decreases in each round shows that this quantity can only decrease.

We next provide necessary and sufficient conditions on the perishing distribution to satisfy the conditions outlined in Theorem 4.1. Since our goal is to highlight the perishing process's dependence on T, for clarity of exposition we assume $B = N = T$ almost surely, and set $\underline{X} = 1$. At the cost of cumbersome algebra, identical insights can be derived when relaxing this assumption. We omit the straightforward proofs of the facts in this section.

The following proposition establishes that, for our conditions to hold, in the worst case at most $O(\sqrt{T})$ items can perish in the first period, in expectation.

Proposition 4.6. *Suppose $\sum_{b\in[B]} \mathbb{P}(T_b = 1) \geq \sqrt{T-1}$. Then, for any ordering σ: (i) Condition* (1) *is not satisfied for any value of δ, and (ii) $\eta_1 > \sqrt{T-1}$.*

Note that this necessary condition fails to hold for one of the most standard models of perishing: geometrically distributed perishing with parameter $1/T$ (that is, a constant fraction of items perish every day). This highlights that one of the most popular models in the literature is, in a sense, far too pessimistic; for this setting, there is no hope of achieving low envy and efficiency.

Having established this necessary condition, we next consider the two extreme models of perishing that have been considered in the literature: i.i.d. and deterministic perishing times. In a sense, the former setting is the most difficult case for the decision-maker, as the best they can do is to choose which items to allocate uniformly at random, resulting in a high number of 'mistakes' in hindsight. In this case, then, it must be that the vast majority of items perish toward the end of the time horizon. We formalize this intuition below.

Proposition 4.7. *Suppose items' perishing times are drawn i.i.d. from the same distribution in each period, and satisfies for all $b \in [B]$, $t \in [T-1]$:*

1. $\mathbb{P}(t \leq T_b < T) \leq \sqrt{\frac{T-t}{(T-t+1)^2}}$
2. $\mathbb{P}(T_b = 1) \geq 1 - \frac{t}{T} - \sqrt{\frac{T-t}{(T-t+1)^2}}$.

Furthermore, let σ allocate items arbitrarily. Then, (i) there exists $\delta^* > 0$ such that Condition (1) *is satisfied $\forall\ \delta \leq \delta^*$, and (ii) $\eta_t \leq \sqrt{T-t}$ for all $t \in [T]$.*

Deterministic perishing times find themselves at the opposite end of the spectrum: obtaining a "good" perishing prediction comes for free. Fairness, however, does not; to see this, suppose all items perish at the end of the first period. At a high level, then, fairness in the deterministic setting is hindered by "redundancy" in the initial budget (i.e. resources that will perish at the same timestep). Proposition 4.8 formalizes this idea.

Proposition 4.8. *Suppose items' perishing times are known, deterministic, and offset-expiring. Let $P_t = \sum_{b\in[B]} \mathbb{1}\{T_b = t\}$ be the number of resources that perish at time t, and suppose $\sum_{\tau \geq t} P_\tau \leq \sqrt{T-t}$ for all t. If σ allocates items in increasing order of T_b, breaking ties arbitrarily, then $\eta_t \leq \sqrt{T-t}$ for all $t \in [T]$.*

In the previous example, the hindrance arises from the fact that the algorithm is unable to allocate resources fast enough relative to the time they are perishing. We close with an example combining these two aspects, both "redundancy" (with multiple resources with identical expected perishing times) and classes of i.i.d. distributions.

Proposition 4.9. *Suppose items are partitioned into two classes of size $\sqrt{T}$ and $T - \sqrt{T}$ respectively. Suppose that within each class the items are drawn i.i.d. from the same distribution, and that the following hold:*

1. $\mathbb{P}(T_b^{(1)} < T) = \frac{1}{2}\sqrt{1 - \frac{1}{\sqrt{T}}}$ *for all items in the first class;*
2. $\mathbb{P}(T_b^{(2)} < T) = \frac{1}{2}\sqrt{\frac{1}{T-\sqrt{T}}}$ *for all items in the second class;*
3. $\sqrt{T}\mathbb{P}(T_b^{(1)} < t) + (T - \sqrt{T})\mathbb{P}(T_b^{(2)} < t) < t$ *for all $t \in [T-1]$.*

Let σ allocate all items in the first class before items in the second class, breaking ties arbitrarily within classes; then (i) there exists $\delta^* > 0$ such that Condition (1) is satisfied $\forall \delta \leq \delta^*$, and (ii) $\eta_t \leq \sqrt{T-t}$ for all $t \in [T]$.

We leave a further investigation of the necessity of the condition on η_t to future work.

5 Online Fair Allocation with Non-Linear Utility Functions

In this section we discuss the extension of our results to a broad class of utility functions, moving past the additive utility assumption typically made in the literature. In particular, we consider the following class of utilities:

Assumption 1. *For $\theta \in \Theta$, $u(\cdot, \theta)$ is concave, one-positively homothetic[2], and strictly increasing.*

We refer to this class of functions as *regular* utilities. Examples include:

- **Linear**: $u(x, \theta) = \langle w_\theta, x \rangle$ for some vector of per-resource preferences $w_\theta \in \mathbb{R}^K_{>0}$.
- **Cobb-Douglas**: $u(x, \theta) = \prod_{k=1}^K w_{\theta,k} x_k^{\alpha_{\theta,k}}$, where $w_\theta \in \mathbb{R}^K_{>0}, \alpha_\theta \in \mathbb{R}^K_{>0}$, and $\sum_{k \in [K]} \alpha_{\theta,k} = 1$.
- **Leontief + ϵ-linear**: Such an ϵ-perturbation trick is more widely useful for utilities which are homothetic and concave, but not strictly increasing.: $u(x, \theta) = \min\{x_k / w_{\theta,k}\} + \epsilon \langle w'_\theta, x_\theta \rangle$, where $w_\theta, w'_\theta \in \mathbb{R}^K_{>0}$, and $\epsilon > 0$.

Note that the popular class of utility functions, Leontief utilities, fails to satisfy the requirements of Proposition 5.2 since they are not *strictly* increasing. In Appendix E we however show that our algorithmic framework can nonetheless be leveraged to obtain fairness bounds in this latter setting by considering the Leontief+ϵ-linear perturbation, for an appropriately chosen value of ϵ.

In order to disentangle the complexity of non-additive utilities with that of the uncertainty of the perishing process, in the remainder of this section we consider the no-perishing setting. Tacking perishing onto these insights then follows immediately via a straightforward modification of the definition of offset-expiring processes (for which the hindsight fair solution is the same as that under no-perishing), and by maintaining separate allocation schedules for each of the K resources. Finally, before proceeding, we emphasize that our contribution here is technical, rather than algorithmic. The algorithmic idea presented in Algorithm 2 was first introduced in Sinclair et al. (2022); their analysis, however, relied heavily on the additive utility assumption. Via careful analysis, we show that the insights continue to hold for this much broader, practical class of utilities.

5.1 Structure of Offline Eisenberg-Gale Solutions

We being by establishing properties of solutions to the EG program under regular utilities that will be leveraged to establish our algorithm's guarantees. We defer their proofs to Appendix D.1.

Recall, one of the fairness metrics we consider is that of *counterfactual envy*, the distance between the utilities under the algorithm's allocations and the hindsight optimal solution produced by the EG program (easily shown to be fair under additive utilities). We first establish that this latter solution remains a meaningful benchmark for regular utilities (i.e., it is indeed fair in hindsight).

The following property of solutions to the EG program will simplify subsequent analyses.

[2] $u(\cdot, \theta)$ is one-positively homothetic if $u(\alpha x, \theta) = \alpha u(x, \theta)$ for all $\alpha > 0$.

Proposition 5.1. *Suppose $u(x,\theta)$ is concave in x for all $\theta \in \Theta$. Then there exists a time-invariant solution X^{opt} to the EG program. That is, for all $\theta \in \Theta$, $X^{opt}_{t,\theta} = X^{opt}_{t',\theta}$ for all $t, t' \in [T]$.*

Thus, in the remainder of the section we will equivalently consider the following formulation of the EG program:

$$\begin{aligned} \max_{X \in \mathbb{R}^{|\Theta| \times K}} \sum_{\theta \in \Theta} & N_\theta \log u(X_\theta, \theta) \\ \text{s.t.} \sum_{\theta \in \Theta} & N_\theta X_{\theta,k} \leq B_k \quad \forall\, k \in [K] \\ & X_{\theta,k} \geq 0 \quad \forall\, \theta \in \Theta, k \in [K]. \end{aligned} \tag{EG}$$

With this simplification in hand, we show that solutions to the EG program satisfy the three fairness properties defined in Definition 3.1.

Proposition 5.2. *Suppose $u(\cdot, \theta)$ satisfies Assumption 1. Then, any solution to the EG program is Pareto-efficient, envy-free, and proportional.*

The proof of Proposition 5.2 relies on the fact that, under the regularity assumption, solutions to the EG program correspond precisely to allocations in a competitive equilibrium of the corresponding Fisher market with equal incomes. We then leverage the seminal result of Varian (1974) which states that, under mild conditions, competitive equilibria are envy-free and Pareto efficient. With some additional work, proportionality follows.

Proposition 5.3, which describes sensitivity of solutions to the EG program, is central to establishing our algorithm's performance guarantee, and is our key technical contribution of this section. To the best of our knowledge the second property, termed *competitive monotonicity* in the literature, was only known to hold for the setting of weak gross substitutes (Jain and Vazirani, 2010).

Proposition 5.3. *Let $x((N_\theta)_{\theta \in \Theta})$ denote optimal primal solutions to the Eisenberg-Gale program for a given vector of arrivals $(N_\theta)_{\theta \in \Theta}$, and fix $\zeta > 0$. Then, we have that:*

(1) *Scaling: If $\widetilde{N}_\theta = (1+\zeta) N_\theta$ for every $\theta \in \Theta$, then:*

$$\begin{aligned} x((\widetilde{N}_\theta)_{\theta \in \Theta}) &= \frac{x((N_\theta)_{\theta \in \Theta})}{1+\zeta} \\ u(x((N_\theta)_{\theta \in \Theta})_\theta, \theta) - u(x((\widetilde{N}_\theta)_{\theta \in \Theta})_\theta, \theta) &= \left(1 - \frac{1}{1+\zeta}\right) u(x((N_\theta)_{\theta \in \Theta})_\theta, \theta) \end{aligned}$$

(2) *Monotonicity: If $N_\theta \leq \widetilde{N}_\theta$ for every $\theta \in \Theta$ then*

$$u(x((\widetilde{N}_\theta)_{\theta \in \Theta})_\theta, \theta) \leq u(x((N_\theta)_{\theta \in \Theta})_\theta, \theta) \quad \forall \theta \in \Theta$$

5.2 Algorithm and Performance Guarantee

Whereas for the setting with additive utilities and $K = 1$, it suffices to construct guardrails such that $\|\underline{X} - \overline{X}\| = L_T$, for more general utility functions an L_T-distance in allocations need not translate into an L_T distance in utilities. Moreover, a key challenge in achieving lower counterfactual envy is the construction of guardrails such that the *utilities* under these allocations sandwiches the utilities under X^{opt}. Our algorithm achieves this by solving the EG program twice: once assuming a high-probability upper bound on the arrival sequence, and once assuming a high-probability lower bound

ALGORITHM 2: NL-GUARDRAIL

Input: Budget $B = B_1^{alg}$, confidence terms $(\textsc{Conf}_{t,t',\theta})_{t,t' \in [T], \theta \in \Theta}$, desired bound on envy L_T, arrival confidence bounds $(\overline{n}_\theta)_{\theta \in \Theta}$, $(\underline{n}_\theta)_{\theta \in \Theta}$

Output: An allocation $X^{alg} \in \mathbb{R}^{T \times |\Theta| \times K}$

Solve for $\overline{X} = x((\underline{n}_\theta)_{\theta \in \Theta})$ as the solution to (EG) with arrival vector $(\underline{n}_\theta)_{\theta \in \Theta}$

Solve for $\underline{X} = x((\overline{n}_\theta)_{\theta \in \Theta})$ as the solution to (EG) with arrival vector $(\overline{n}_\theta)_{\theta \in \Theta}$ `// solve for guardrails`

for $t = 1, \ldots, T$ **do**

  for *each resource* $k \in [K]$ **do**

    if $B_{t,k}^{alg} < \sum_{\theta \in \Theta} N_{t,\theta} \underline{X}_{\theta,k}$ **then** `// insufficient budget to allocate lower guardrail`

      Set $X_{t,\theta,k}^{alg} = \frac{B_{t,k}^{alg}}{\sum_{\theta \in \Theta} N_{t,\theta}}$ for each $\theta \in \Theta$

    else if $B_{t,k}^{alg} - \sum_\theta N_{t,\theta} \overline{X}_{\theta,k} \geq \sum_{\theta \in \Theta} \underline{X}_{\theta,k} (\mathbb{E}[N_{>t,\theta}] + \textsc{Conf}_{t,T,\theta})$ **then**`// use upper guardrail`

      Set $X_{t,\theta,k}^{alg} = \overline{X}_{\theta,k}$ for each $\theta \in \Theta$

    else `// use lower guardrail`

      Set $X_{t,\theta,k}^{alg} = \underline{X}_{\theta,k}$ for each $\theta \in \Theta$

  end

  Update $B_{t+1}^{alg} = B_t^{alg} - \sum_{\theta \in \Theta} N_{t,\theta} X_{t,\theta}^{alg}$

end

return X^{alg}

on the arrival sequence. It then allocates either of the two guardrails similarly to Algorithm 1. Our algorithm, NL-GUARDRAIL, is presented in Algorithm 2.

Our main result shows that, under carefully constructed guardrails, NL-GUARDRAIL achieves the optimal envy-efficiency trade-off for all regular utility functions.

Theorem 5.4. *Fix $L_T = o(1)$, and let* NL-GUARDRAIL *be initialized with*

$$\begin{aligned}
\textsc{Conf}_{t,t',\theta} &= \sqrt{2(t'-t)\rho_{max}^2 \log(2T^2|\Theta|/\delta)} \quad \forall\, t' > t, \theta \in \Theta \\
\overline{n}_\theta &= \mathbb{E}[N_\theta]\left(1 + \max_\theta \frac{\textsc{Conf}_{0,T,\theta}}{\mathbb{E}[N_\theta]}\right) \quad \forall\, \theta \in \Theta \\
\underline{n}_\theta &= \mathbb{E}[N_\theta](1-c) \quad \forall\, \theta \in \Theta,\ c = \frac{L_T}{\max_{\theta \in \Theta} u(B,\theta)}\left(1 + \max_\theta \frac{\textsc{Conf}_{0,T,\theta}}{\mathbb{E}[N_\theta]}\right) - \max_\theta \frac{\textsc{Conf}_{0,T,\theta}}{\mathbb{E}[N_\theta]}.
\end{aligned}$$

Then, with probability at least $1-\delta$, NL-GUARDRAIL *achieves:*

$$\textsc{Envy} \leq L_T \qquad \Delta_{EF} \lesssim \max\{1/\sqrt{T}, L_T\} \qquad \Delta_{\mathit{efficiency}} \lesssim \min\{\sqrt{T}, 1/L_T\} \qquad \Delta_{prop} \lesssim \max\{1/\sqrt{T}, L_T\},$$

where $\lesssim$ drops poly-logarithmic factors of T, $\log(1/\delta)$, $o(1)$ terms, and absolute constants.

The outline of the proof follows similarly to that of Theorem 4.1. Namely, we restrict our attention to a "good event" $\mathcal{E} = \{|N_{(t,t'],\theta} - \mathbb{E}\big[N_{(t,t'],\theta}\big]| \leq \textsc{Conf}_{t,t',\theta}\ \forall\ t' > t, \theta \in \Theta\}$ which holds with probability $1-\delta$. We then show that, given $\mathcal{E}$, the algorithm never runs out budget, and is adaptively cautious. The main obstacle in establishing our bounds lies in (i) analyzing agents' utilities under $\overline{X}$ and $\underline{X}$, static solutions to the EG program which are fair, and (ii) relating these back to X^{opt}, the offline optimal solution, all as a function of L_T. Proposition 5.3 is central in tackling this challenge. We defer the proof of the theorem to Appendix D.2.

6 Numerical Experiments

We conclude by complementing the theoretical analysis of Perishing-Guardrail (Section 4) and NL-Guardrail (Section 5) with a numerical study using data provided by a food bank adapted from Sinclair et al. (2022). For each of the models described in Sections 4 and 5, we first describe the data-driven experiments, and then compare the effectiveness of our algorithms to that of state-of-the-art benchmarks. All of the code for the experiments is available at `https://github.com/seanrsinclair/Online-Resource-Allocation`. Due to space limitations, see Appendix F for experiments with non-linear utilities.

6.1 Experimental setup

In each simulation we set the total budget B to be $\sum_{t,\theta} \mathbb{E}[N_{t,\theta}]$, so that the average budget per unit demand is $\beta_{avg} = 1$.

For our first set of experiments, we consider a single type of perishable resource with individual utilities $u(x, \theta) = x$, for simplicity. We test the performance of our algorithm for three different perishing distributions satisfying the conditions outlined in Section 4.3 (motivated by the examples in Section 4.3):

1. **I.I.D.**: $T_b \in \{T/2, T\}$, with $\mathbb{P}(T_b = T/2) = \sqrt{\frac{1}{2}\frac{T}{(\frac{1}{2}T+1)^2}}$.
2. **Deterministic**: Deterministic perishing times satisfying offset-expiry are constructed such that the number of duplicate items perishing after t is exactly $\sqrt{T-t}$.
3. **Two-Class**: $T_b \in \{T/2, T\}$, with the first $\sqrt{T}$ resources perishing according to $\mathbb{P}(T_b^{(1)} = T/2) = \frac{1}{2}\sqrt{1 - \frac{1}{\sqrt{T}}}$, and the remaining resources perishing according to $\mathbb{P}(T_b^{(2)} = T/2) = \frac{1}{2}\sqrt{\frac{1}{T-\sqrt{T}}}$.

6.2 Simulation Results

For each of the experiments, we compare the performance of Perishing-Guardrail on values of $L_T \in \{0, T^{-1/2}, T^{-1/3}\}$ to two state-of-the-art algorithms:

1. Guarded-Hope (Sinclair et al. (2022)) which similarly takes as input an envy parameter L_T, but does not account for perishing;
2. Static CE, which solves (EG-P), replacing the random variables $(N_t, P_t)_{t=1}^T$ with their expectation, and allocating according to the initial solution until it runs out of budget.

Our results can be found in Fig. 1.

In all settings, running Perishing-Guardrail with $L_T > 0$ vastly outperforms the zero-envy variant, with $L_T = 0$. This highlights how the use of both upper and lower guardrails improves efficiency with little-to-no trade-off on envy. We moreover observe relatively little empirical difference in performance on $\Delta_{\textit{efficiency}}$ between Perishing-Guardrail for $L_T = T^{-1/2}$ and $T^{-1/3}$; however, Δ_{EF} is much smaller for $L_T = T^{-1/2}$. This suggests that setting $L_T = T^{-1/2}$ achieves strong performance on Δ_{EF} with relatively little trade-off in $\Delta_{\textit{efficiency}}$.

We also observe that Perishing-Guardrail outperforms Guarded-Hope across all settings (with the exception of the **Deterministic** setup, for which perishing plays no role due to its slow and non-random nature). This highlights how accounting for future perishing is necessary for designing

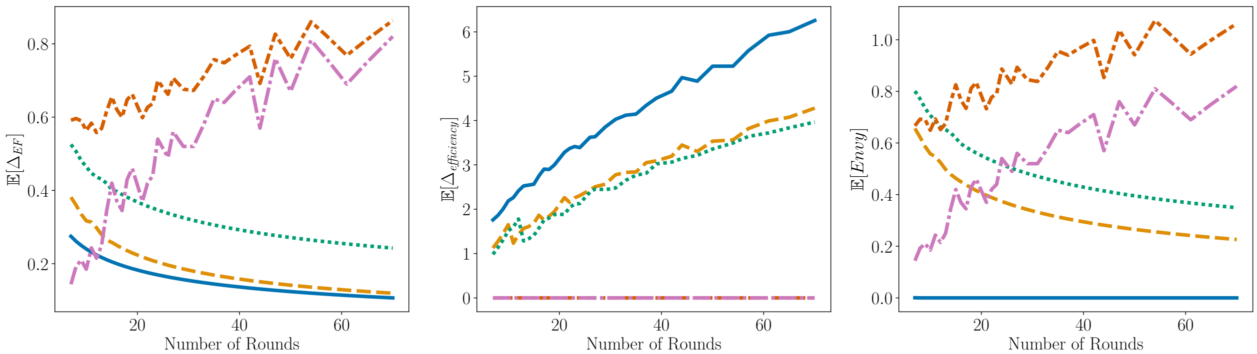

(a) **I.I.D.** (Proposition 4.7)

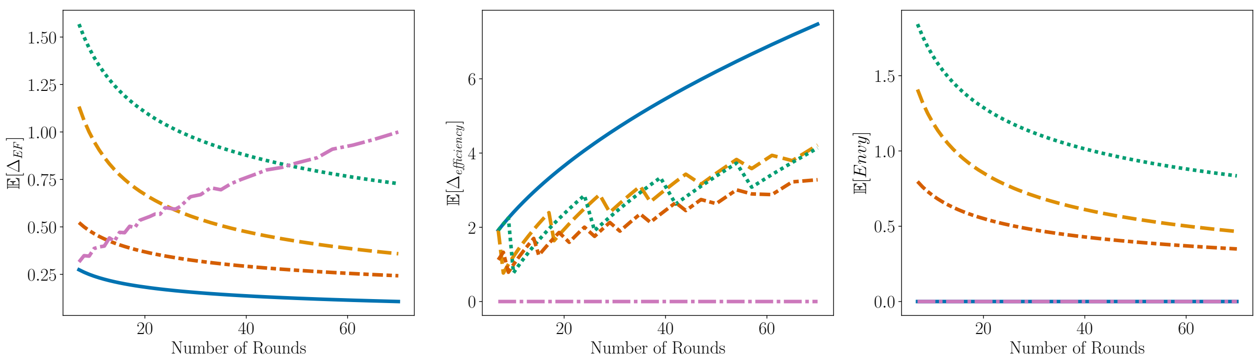

(b) **Deterministic** (Proposition 4.8)

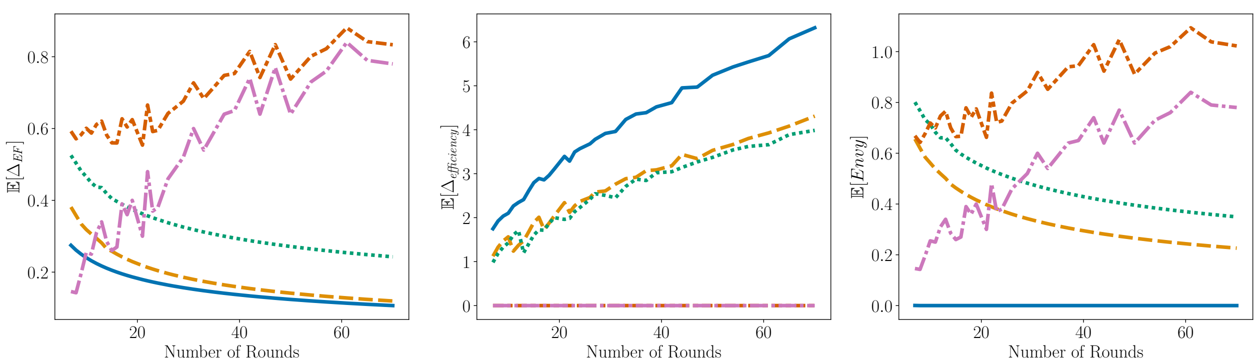

(c) **Two Classes** (Proposition 4.9)

Figure 1: *Comparison of* (*i*) Perishing-Guardrail *for* $L_T = 0, T^{-1/2}$ *and* $T^{-1/3}$, (*ii*) Guarded-Hope *from Sinclair et al. (2022) with* $L_T = T^{-1/2}$, *and* (*iii*) Static CE.

an algorithm: without it, the algorithm fails to allocate resources that are set to perish, as a result running out and yielding poor performance on Δ_{EF}. Finally, we see that Static CE achieves poor performance on Δ_{EF}. Such behavior is intuitive since with constant probability arrivals come in higher than expected; as a result the EG program's solution allocates too aggressively and similarly runs out of budget.

7 Conclusion

In this paper we considered a practically motivated variant of the canonical problem of online fair allocation: a principal has a budget B of *perishable* resources to allocate over T rounds. Each round sees a random number of arrivals with potentially *non-linear utility functions*, and the principal must commit to an allocation before moving on to the next round. We derived an algorithm which, under "good" allocation schedules, and "regular" utilities, achieves the optimal envy-efficiency trade-off under *no perishing*. In so doing, we demonstrate the effectiveness of guardrail allocations for a much wider variety of settings than what was initially known.

Several open questions remain, including (*i*) investigating the necessity of the condition on η_t to achieve this trade-off, (*ii*) deriving the envy-efficiency Pareto frontier in relation to *total* waste (include spoilage), and (*iii*) considering more appropriate notions of envy for non-offset-expiring processes. Finally, though this paper considered exogenous *depletion* of the budget, a natural practical extension is one wherein B evolves stochastically, accounting for external donations independent of the allocations made by the algorithm.

Acknowledgments

Part of this work was done while Sean Sinclair and Sid Banerjee were visiting the Simons Institute for the Theory of Computing for the semester on Data-Driven Decision Processes. We gratefully acknowledge funding from the National Science Foundation under grants ECCS-1847393, DMS-1839346, CCF-1948256, CNS-195599, and CNS-1955997, the Air Force Office of Scientific Research under grant FA9550-23-1-0068, and the Army Research Laboratory under grants W911NF-19-1-0217 and W911NF-17-1-0094. We would also like to thank the Food Bank of the Southern Tier and the Cornell Mathematical Contest in Modeling for their collaboration.

References

Aleksandrov, M., Aziz, H., Gaspers, S., and Walsh, T. (2015). Online fair division: Analysing a food bank problem. In *Proceedings of the 24th International Conference on Artificial Intelligence*, IJCAI'15, page 2540–2546. AAAI Press.

Aleksandrov, M. and Walsh, T. (2019a). Monotone and online fair division. In *Joint German/Austrian Conference on Artificial Intelligence (Künstliche Intelligenz)*, pages 60–75. Springer.

Aleksandrov, M. and Walsh, T. (2019b). Online fair division: A survey. *arXiv preprint arXiv:1911.09488*.

Alijani, R., Banerjee, S., Gollapudi, S., Munagala, K., and Wang, K. (2020). Predict and match: Prophet inequalities with uncertain supply. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(1):1–23.

Aouad, A. and Saritaç, Ö. (2020). Dynamic stochastic matching under limited time. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 789–790.

Aristotle Cloud Federation Project (2022). https://www.hpcwire.com/off-the-wire/aristotle-project-advances-campus-cloud-technologies/.

Azar, Y., Buchbinder, N., and Jain, K. (2010). How to allocate goods in an online market? In *European Symposium on Algorithms*, pages 51–62. Springer.

Aziz, H., Schlotter, I. A., and Walsh, T. (2016). Control of fair division. IJCAI.

Banerjee, S., Gkatzelis, V., Gorokh, A., and Jin, B. (2020). Online nash social welfare maximization via promised utilities. *arXiv preprint arXiv:2008.03564*.

Bansal, N., Jiang, H., Singla, S., and Sinha, M. (2020). *Online Vector Balancing and Geometric Discrepancy*, page 1139–1152. Association for Computing Machinery, New York, NY, USA.

Bateni, M., Chen, Y., Ciocan, D. F., and Mirrokni, V. (2022). Fair resource allocation in a volatile marketplace. *Operations Research*, 70(1):288–308.

Benade, G., Kazachkov, A. M., Procaccia, A. D., and Psomas, C.-A. (2018). How to make envy vanish over time. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 593–610.

Bertsimas, D., Farias, V. F., and Trichakis, N. (2011). The price of fairness. *Operations research*, 59(1):17–31.

Bogomolnaia, A., Moulin, H., and Sandomirskiy, F. (2021). On the fair division of a random object. *Management Science*.

Bonald, T., Massoulié, L., Proutiere, A., and Virtamo, J. (2006). A queueing analysis of max-min fairness, proportional fairness and balanced fairness. *Queueing systems*, 53:65–84.

Cole, R., Gkatzelis, V., and Goel, G. (2013). Mechanism design for fair division: Allocating divisible items without payments. In *EC 2013*.

Donahue, K. and Kleinberg, J. (2020). Fairness and utilization in allocating resources with uncertain demand. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 658–668.

Eaton, J. (2022). The u.s. has wasted over 82 million covid vaccine doses. *NBC News*.

Eisenberg, E. (1961). Aggregation of utility functions. *Management Science*, 7(4):337–350.

Elzayn, H., Jabbari, S., Jung, C., Kearns, M., Neel, S., Roth, A., and Schutzman, Z. (2019). Fair algorithms for learning in allocation problems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 170–179.

Freeman, R., Zahedi, S. M., and Conitzer, V. (2017). Fair social choice in dynamic settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4580–4587. International Joint Conferences on Artificial Intelligence Marina del Rey, CA.

Freund, D. and Hssaine, C. (2021). Fair incentives for repeated engagement. *arXiv preprint arXiv:2111.00002*.

Friedman, E., Psomas, C.-A., and Vardi, S. (2015). Dynamic fair division with minimal disruptions. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, page 697–713, New York, NY, USA. Association for Computing Machinery.

Friedman, E., Psomas, C.-A., and Vardi, S. (2017). Controlled dynamic fair division. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, page 461–478, New York, NY, USA. Association for Computing Machinery.

Gerding, E., Perez-Diaz, A., Aziz, H., Gaspers, S., Marcu, A., Mattei, N., and Walsh, T. (2019). Fair online allocation of perishable goods and its application to electric vehicle charging.

Ghodsi, A., Zaharia, M., Hindman, B., Konwinski, A., Shenker, S., and Stoica, I. (2011). Dominant resource fairness: Fair allocation of multiple resource types. In *Nsdi*, volume 11, pages 24–24.

Grosof, I., Scully, Z., and Harchol-Balter, M. (2019). Load balancing guardrails: keeping your heavy traffic on the road to low response times. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 3(2):1–31.

Hanukov, G., Avinadav, T., Chernonog, T., and Yechiali, U. (2020). A service system with perishable products where customers are either fastidious or strategic. *International Journal of Production Economics*, 228:107696.

Hassanzadeh, P., Kreacic, E., Zeng, S., Xiao, Y., and Ganesh, S. (2023). Sequential fair resource allocation under a markov decision process framework. *arXiv preprint arXiv:2301.03758*.

He, J., Procaccia, A. D., Psomas, A., and Zeng, D. (2019). Achieving a fairer future by changing the past. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 343–349. International Joint Conferences on Artificial Intelligence Organization.

Jain, K. and Vazirani, V. V. (2010). Eisenberg–gale markets: algorithms and game-theoretic properties. *Games and Economic Behavior*, 70(1):84–106.

Joe-Wong, C., Sen, S., Lan, T., and Chiang, M. (2013). Multiresource allocation: Fairness–efficiency tradeoffs in a unifying framework. *IEEE/ACM Transactions on Networking*, 21(6):1785–1798.

Kalinowski, T., Narodytska, N., and Walsh, T. (2013). A social welfare optimal sequential allocation procedure. In *Twenty-Third International Joint Conference on Artificial Intelligence*.

Kelly, F. P., Maulloo, A. K., and Tan, D. K. H. (1998). Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research society*, 49(3):237–252.

Lien, R. W., Iravani, S. M., and Smilowitz, K. R. (2014). Sequential resource allocation for nonprofit operations. *Operations Research*, 62(2):301–317.

Lupkin, S. (2020). How feds decide on remdesivir shipments to states remains mysterious.

Manshadi, V., Niazadeh, R., and Rodilitz, S. (2021). Fair dynamic rationing. *Available at SSRN 3775895*.

Mattei, N., Saffidine, A., and Walsh, T. (2017). Mechanisms for online organ matching. In *IJCAI*, pages 345–351.

Mattei, N., Saffidine, A., and Walsh, T. (2018). Fairness in deceased organ matching. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 236–242.

Nahmias, S. (2011). *Perishable inventory systems*, volume 160. Springer Science & Business Media.

Nisan, N., Roughgarden, T., Tardos, É., and Vazirani, V. V., editors (2007). *Algorithmic Game Theory*. Cambridge University Press.

Perry, D. (1999). Analysis of a sampling control scheme for a perishable inventory system. *Operations Research*, 47(6):966–973.

Prendergast, C. (2017). How food banks use markets to feed the poor. *Journal of Economic Perspectives*, 31(4):145–62.

Sinclair, S. R., Jain, G., Banerjee, S., and Yu, C. L. (2022). Sequential fair allocation: Achieving the optimal envy-efficiency trade-off curve. *Operations Research*.

Varian, H. R. (1974). Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91.

Varian, H. R. (1976). Two problems in the theory of fairness. *Journal of Public Economics*, 5(3-4):249–260.

Vershynin, R. (2018). *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press.

Zeng, D. and Psomas, A. (2020). Fairness-efficiency tradeoffs in dynamic fair division. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20, page 911–912, New York, NY, USA. Association for Computing Machinery.

A Table of Notation

Symbol	Definition
	Problem setting specifications
T	Total number of rounds
K, B_k	Number of resources and budget for resource $k \in [K]$
Θ, θ	Set of types for individuals, and specification for individual's type
$u(x, \theta) : \mathbb{R}^k \times \Theta \to \mathbb{R}_+$	Utility function for individuals of type θ
$N_{t,\theta}$	Number of individuals of type θ in round t
$\mathcal{F}_t$	Known distribution over $(N_{t,\theta})_{\theta\in\Theta}$
$N_{\geq t,\theta}$	$\sum_{t' \geq t} N_{t',\theta}$
$\sigma_{t,\theta}, \rho_{t,\theta}, \mu_{t,\theta}$	$\mathrm{Var}[N_{t,\theta}]$, bound on $\lvert N_{t,\theta} - \mathbb{E}[N_{t,\theta}]\rvert$, and $\mathbb{E}[N_{t,\theta}]$
$\sigma^2_{min}, \sigma^2_{max}$	The respective maximum and minimum value of each quantity
T_b, P_t	Perishing time for resource $b \in [B]$ and $P_t = \sum_b \mathbb{1}\{T_b = t\}$
β_{avg}	$\sum_k B_k / \sum_{\theta\in\Theta} \mathbb{E}[N_\theta]$
X^{opt}, X^{alg}	Optimal fair allocation in hindsight and allocation by algorithm
Δ_{EF}	$\max_{t\in[T],\theta\in\Theta} \lVert u(X^{alg}_{t,\theta}, \theta) - u(X^{opt}_{t,\theta}, \theta)\rVert_\infty$
Envy	$\max_{t,t'\in[T]^2,\theta,\theta'\in\Theta^2} u(X^{alg}_{t',\theta'}, \theta) - u(X^{alg}_{t,\theta}, \theta)$
$\Delta_{\mathit{efficiency}}$	$\sum_k B_k - \sum_{\theta\in\Theta} \sum_t N_{t,\theta} X^{alg}_{t,\theta,k} - \sum_{t,k} \mathrm{PUA}_{t,k}$
	Algorithm specification
B^{alg}_t	Budget available to the algorithm at start of round t
Conf_t	Confidence bound on $N_{t,\theta}$ and P_t, indicated by superscript
L_T	Desired bound on Δ_{EF}, Envy
σ	Pre-specified allocation schedule
δ	High probability constant
	Additional notation
$\Phi(\cdot)$	Standard normal CDF
$\bar{\rho}_\theta, \bar{\sigma}^2_\theta, \bar{\mu}_\theta$	Averages of these quantities, i.e. $\frac{1}{T}\sum_t \rho_{t,\theta}$, $\frac{1}{T}\sum_t \sigma^2_{t,\theta}$, $\frac{1}{T}\sum_t \mathbb{E}[N_{t,\theta}]$
$\mathcal{A}_t$	Set of resources allocated by algorithm in round t
PUA_t	Resources which perish and are un-allocated in round t,$\sum_b \mathbb{1}\{T_b = t, b \notin \mathcal{A}_t\}$
$\mathrm{Prec}(b \mid S)$	$\sum_{b'\in S} \mathbb{1}\{b' \preceq_\sigma b\}$
τ_b	Period in which b would have been allocated after t under slow consumption
	$\tau_b(t \mid S) = \inf\{t' \geq t : \underline{N}_{[t:t']}\underline{X} > \mathrm{Prec}(b \mid S)\}$.
$[\overline{B}_t]$	$\{\sigma(\lfloor \underline{N}_{<t}\underline{X} \rfloor), \ldots, \sigma(B)\}$
η_t	$\sum_{b\in[\overline{B}_t]} \mathbb{P}\big(t \leq T_b < \min\{T, \tau_b(t \mid [\overline{B}_t])\}\big)$

Table 1: Common notation

B Section 3 Proofs

Proof. Let $(p_t)_{t\in[T]}$ and $(\alpha_{t,\theta})_{t\in[T],\theta\in\Theta}$ denote the dual variables corresponding to the budget constraints and nonnegativity constraints, respectively. Then, by the KKT conditions, we have:

1. $p_t > 0 \implies \sum_{\tau\geq t}\sum_\theta N_{\tau,\theta}X_{\tau,\theta} \leq B - P_{<t}$
2. $\alpha_{t,\theta} > 0 \implies X_{t,\theta} = 0$
3. $\frac{1}{X_{t,\theta}} \leq \sum_{\tau\leq t} p_\tau$, with equality whenever $X_{t,\theta} > 0$.

Let $p_1 = N/B$, and $p_t = 0$ for all $t \geq 2$. This construction satisfies properties (1) and (3). It thus suffices to check feasibility, i.e. that $\sum_{\tau\geq t}\sum_\theta N_{\tau,\theta}X_{\tau,\theta} \leq B - P_{<t}$ for all t, for $X_{t,\theta} = B/N$. This holds if, for all t:

$$\frac{B}{N}N_{\geq t} \leq B - P_{<t} \iff P_{<t} \leq B\left(1 - \frac{N_{\geq t}}{N}\right) = B\cdot\frac{N_{<t}}{N},$$

which holds by assumption. □

C Section 4 Proofs

C.1 Theorem 4.1 auxiliary lemmas

C.1.1 Proof of Lemma 4.2

The result is a corollary of high-probability upper bounds on $\mathcal{E}_N$, $\mathcal{E}_P$ and $\mathcal{E}_{\overline{P}}$.

Lemma C.1 (Concentration on Arrival Process)**.** *$\mathcal{E}_N$ holds with probability at least $1-\delta$.*

Proof. Fix $t' < t$. Recall, for all $\tau \in [T]$, $\rho_{\tau,\theta} := |N_{\tau,\theta} - \mathbb{E}[N_{\tau,\theta}]|$, which implies

$$N_{\tau,\theta} \in [\mathbb{E}[N_{\tau,\theta}] - \rho_{\tau,\theta}, \mathbb{E}[N_{\tau,\theta}] + \rho_{\tau,\theta}].$$

Thus, from a simple application of Hoeffding's inequality (Theorem H.1):

$$\mathbb{P}(|N_{(t,t']} - \mathbb{E}\big[N_{(t,t']}\big]| \geq \epsilon) \leq 2\exp\left(-\frac{2\epsilon^2}{\sum_\theta\sum_{\tau\in(t,t']}4\rho^2_{\tau,\theta}}\right) \tag{5}$$

We now consider our desired bound.

$$\begin{aligned}\mathbb{P}(|N_{(t,t']} - \mathbb{E}\big[N_{(t,t']}\big] \leq \epsilon\,\forall\, t,t') &\geq 1 - \sum_{t,t'}\mathbb{P}(|N_{(t,t']} - \mathbb{E}\big[N_{(t,t']}\big]| \geq \epsilon) \\ &\geq 1 - \sum_{t,t'} 2\exp\left(-\frac{2\epsilon^2}{\sum_\theta\sum_{\tau\in(t,t']}4\rho^2_{\tau,\theta}}\right) \\ &\geq 1 - \sum_{t,t'} 2\exp\left(-\frac{\epsilon^2}{2|\Theta|\rho^2_{\max}(t'-t)}\right)\end{aligned}$$

where the first inequality follows from a union bound, the second inequality by plugging in Hoeffding's bound (5), and the third inequality by upper bounding $\rho_{\tau,\theta}$ by $\rho_{\max}$, for all $\tau \in (t,t']$. Solving for ϵ such that $2\exp\left(-\frac{\epsilon^2}{2|\Theta|\rho^2_{\max}(t'-t)}\right) = \delta/T^2$, we obtain our result. □

Lemma C.2 (Concentration on Perishing Process)**.** *$\mathcal{E}_P$ holds with probability at least $1-\delta$.*

Proof. By definition, $P_{<t} = \sum_{b\in[B]} \mathbb{1}\{T_b < t\}$, and $\mathbb{E}[P_{<t}] = \sum_{b\in[B]} \mathbb{P}(T_b < t)$. Moreover, for $b \in [B]$, $\mathbb{1}\{T_b < t\}$ is a Bernoulli random variable with probability of success $\mathbb{P}(T_b < t)$. Using Theorem H.2 we then have that:

$$\mathbb{P}(|P_{<t} - \nu_t| \geq \epsilon\nu_t) \leq 2\exp\left(-\frac{\nu_t\epsilon^2}{3}\right).$$

Setting the right-hand side equal to δ/T and solving for ϵ yields $\epsilon = \sqrt{\frac{3}{\nu_t}\log(2T/\delta)}$. Thus we have that with probability at least $1-\delta/T$:

$$|P_{<t} - \nu_t| \leq \nu_t\epsilon = \sqrt{3\nu_t\log(2T/\delta)}.$$

Taking a union bound over t yields the desired result. □

Lemma C.3. *Given $\mathcal{E}_N$, $\mathcal{E}_{\overline{P}}$ holds with probability at least $1-\delta$.*

Proof. Given B_t^{alg}, we first upper bound $\text{PUA}_{\geq t}$ as a function of the worst-case perishing times $\tau_b(t \mid [\overline{B}_t])$. Recall, $\text{PUA}_{\geq t}$ represents the amount of unallocated goods that perished between t and $T-1$. Formally:

$$\begin{aligned}
\text{PUA}_{\geq t} &\leq \sum_{\tau=t}^{T-1} \sum_{b\in[B_\tau^{alg}]} \mathbb{1}\{T_b = \tau, b \notin \mathcal{A}_\tau\} \\
&\leq \sum_{b\in[B_t^{alg}]} \mathbb{1}\Big\{t \leq T_b < \min\{T, \tau_b(t \mid [B_t^{alg}])\}\Big\}, \qquad (6)
\end{aligned}$$

where (6) follows from the fact that, under $\mathcal{E}_N$, $\underline{N}_{[t:t']} \leq N_{[t:t']}$ for all $t' \in [T]$ and, as long as the algorithm hasn't run out of budget (at which point perishing is no longer possible), the least possible amount that could have been allocated is $\underline{X}$. Thus, $\tau_b(t \mid [B_t^{alg}])$ upper bounds the time at which b would have been allocated under the true arrival sequence $(N_{t',\theta})_{t'\geq t,\theta\in\Theta}$, since it could have been that an item $b' \prec_\sigma b$ perished early, resulting in an earlier allocation date for b.

Via a similar argument, it is easy to see that $[B_t^{alg}] \subseteq [\overline{B}_t]$. Hence, $\textsc{Prec}(b \mid [B_t^{alg}]) \leq \textsc{Prec}(b \mid [\overline{B}_t])$; it then follows that $\tau_b(t \mid [B_t^{alg}]) \leq \tau_b(b \mid [\overline{B}_t])$. Plugging this into (6), we obtain:

$$\text{PUA}_{\geq t} \leq \sum_{b\in[\overline{B}_t]} \mathbb{1}\big\{t \leq T_b < \min\{T, \tau_b(t \mid [\overline{B}_t])\}\big\}. \qquad (7)$$

Applying a Chernoff bound to the sum of independent Bernoulli random variables on the left-hand side of (7) (see Theorem H.2), we obtain that, with probability at least $1-\delta/T$:

$$\begin{aligned}
\text{PUA}_{\geq t} &\leq \sum_{b\in[\overline{B}_t]} \mathbb{P}(t \leq T_b < \min\{T, \tau_b(t \mid [\overline{B}_t])\}) \\
&\quad + \sqrt{3\log(2T/\delta) \sum_{b\in[\overline{B}_t]} \mathbb{P}(t \leq T_b < \min\{T, \tau_b(t \mid [\overline{B}_t])\})} \\
&= \eta_t + \sqrt{3\log(2T/\delta)\eta_t} \\
&= \overline{P}_t,
\end{aligned}$$

by definition. A union bound over t completes the proof of the result. □

Proof. The final high-probability bound follows from straightforward algebra, putting Lemmas C.1, C.2 and C.3 together. ☐

C.1.2 Proof of Lemma 4.3

Proof. By induction on t.
Base Case: $t = 1$.
By construction:

$$\begin{aligned} N\underline{X} &= N \frac{B}{\mathbb{E}[N] + \textsc{Conf}^N_{0,T}} \left(1 - \frac{c}{\sqrt{T}}\right) \\ &\le (\mathbb{E}[N] + \textsc{Conf}^N_{0,T}) \frac{B}{\mathbb{E}[N] + \textsc{Conf}^N_{0,T}} \left(1 - \frac{c}{\sqrt{T}}\right) \quad \text{under event } \mathcal{E} \\ &\le B = B^{alg}_1. \end{aligned}$$

Step Case: $t - 1 \to t$.
In the step case we further condition on the previous allocations made by the algorithm, namely $(X_\tau)_{\tau \le t}$.

Case 1: $X_\tau = \underline{X}$ for all $\tau \le t$.
Since $\underline{X}$ was allocated for all $\tau \le t$, by the recursive budget update:

$$B^{alg}_t = B - N_{<t}\underline{X} - \text{PUA}_{<t},$$

where $\text{PUA}_{<t}$ denotes the quantity of unallocated goods that perished before the end of round t. Re-arranging, $B^{alg}_t \ge N_{\ge t}\underline{X}$ is then equivalent to $B \ge N\underline{X} + \text{PUA}_{<t}$. Since $\text{PUA}_{<t}$ is non-decreasing with respect to t, it suffices to show that this holds for $t = T$, or that:

$$B \ge N\underline{X} + \text{PUA}_{<T}.$$

Noting that $\text{PUA}_{<T} = \text{PUA}_{\ge 1}$, we apply Lemma C.3 and obtain that the right-hand side is upper bounded by $N\underline{X} + \overline{P}_1$. Additionally using the fact that $N \le \overline{N}$ under $\mathcal{E}_N$, we have:

$$\begin{aligned} N\underline{X} + \overline{P}_1 &\le \overline{N}\underline{X} + \overline{P}_1 \\ &= B\left(1 - \frac{c}{\sqrt{T}}\right) + \eta_1 + \sqrt{3\log(2T/\delta)\eta_1} \\ &\le B\left(1 - \frac{c}{\sqrt{T}}\right) + (1 + \sqrt{3\log(2T/\delta)})\sqrt{T}, \end{aligned}$$

where the equality follows by definition of $\underline{X}$, and the final inequality follows from the assumption that $\eta_1 \le \sqrt{T}$. Thus, it suffices to show that

$$B \ge B\left(1 - \frac{c}{\sqrt{T}}\right) + (1 + \sqrt{3\log(2T/\delta)})\sqrt{T},$$

or equivalently that

$$Bc/\sqrt{T} \ge (1 + \sqrt{3\log(2T/\delta)})\sqrt{T}.$$

Recall, $c = (1 + \sqrt{3\log(2T/\delta)})/\beta_{avg} = (1 + \sqrt{3\log(2T/\delta)}) \cdot \mathbb{E}[N]/B$, by definition. Multiplying by $B/\sqrt{T}$, and using the fact that $N_t \ge 1$ almost surely (and thus $\mathbb{E}[N] \ge T$), we obtain the desired inequality.

Case 2: There exists $\tau < t$ such that $X_\tau = \overline{X}$.

Let $t^* = \sup\{\tau < t : X_\tau = \overline{X}\}$. We have:

$$B_t^{alg} = B_{t^*}^{alg} - N_{t^*}\overline{X} - N_{(t^*,t)}\underline{X} - \text{PUA}_{[t^*,t)}.$$

Since $\overline{X}$ was allocated at t^*, by construction it must have been that $B_{t^*}^{alg} \geq N_{t^*}\overline{X} + \overline{N}_{>t^*}\underline{X} + \overline{P}_{t^*}$. Plugging this into the above and simplifying we have:

$$\begin{aligned} B_t^{alg} &\geq N_{t^*}\overline{X} + \overline{N}_{>t^*}\underline{X} + \overline{P}_{t^*} - N_{t^*}\overline{X} - N_{(t^*,t)}\underline{X} - \text{PUA}_{[t^*,t)} \\ &= \overline{N}_{>t^*}\underline{X} - N_{(t^*,t)}\underline{X} + \overline{P}_{t^*} - \text{PUA}_{[t^*,t)} \\ &\geq N_{\geq t}\underline{X} + \overline{P}_{t^*} - \text{PUA}_{[t^*,t)}, \end{aligned}$$

where the second inequality follows from the fact that, under the good event, $\overline{N}_{>t^*} \geq N_{>t^*}$. Thus, it suffices to show that $\overline{P}_{t^*} \geq \text{PUA}_{[t^*,t)}$. This holds since $\text{PUA}_{[t^*,t)} \leq \text{PUA}_{\geq t^*} \leq \overline{P}_{t^*}$ by Lemma C.3. □

Proof. In the remainder of the proof, we assume that $T - \tilde{c}L_T^{-2} > 0$ (the other case holds trivially). Suppose for contradiction that $0 < t_0 < T - \tilde{c}L_T^{-2}$. By definition of t_0, it must be that the algorithm allocated $X_{t_0}^{alg} = \underline{X}$ and $X_t^{alg} = \overline{X}$ for all $t > t_0$. Then, by construction, for all $t > t_0$,

$$\begin{aligned} \overline{N}_{>t}\underline{X} &\leq B_t^{alg} - N_t\overline{X} - \overline{P}_t \\ &= \left(B_{t_0}^{alg} - N_{t_0}\underline{X} - \text{PUA}_{t_0} - \sum_{\tau=t_0+1}^{t-1}\left(N_\tau\overline{X} + \text{PUA}_\tau\right)\right) - N_t\overline{X} - \overline{P}_t, \end{aligned} \tag{8}$$

where the equality follows from the recursive budget update and the fact that, from t_0+1 onwards, $\underline{X}$ was allocated by assumption.

Since $X_{t_0}^{alg} = \underline{X}$, it must have been that $B_{t_0}^{alg} < N_{t_0}\overline{X} + \overline{N}_{>t_0}\underline{X} + \overline{P}_{t_0}$. Plugging this into (8), we have:

$$\begin{aligned} &\overline{N}_{>t}\underline{X} < N_{t_0}\left(\overline{X} - \underline{X}\right) + \overline{N}_{>t_0}\underline{X} + \overline{P}_{t_0} - \text{PUA}_{[t_0,t)} - N_{(t_0,t]}\overline{X} - \overline{P}_t \\ \iff &N_{(t_0,t]}\overline{X} < N_{t_0}\left(\overline{X} - \underline{X}\right) + \left(\overline{N}_{>t_0} - \overline{N}_{>t}\right)\underline{X} + \overline{P}_{t_0} - \overline{P}_t - \text{PUA}_{[t_0,t)} \end{aligned}$$

Rearranging the inequality after plugging in the definitions of $\overline{N}_{>t_0}$ and $\overline{N}_{>t}$, we obtain:

$$\overline{X}(N_{(t_0,t]} - \mathbb{E}\left[N_{(t_0,t]}\right]) < (\overline{X} - \underline{X})(N_{t_0} - \mathbb{E}\left[N_{(t_0,t]}\right]) + \underline{X}(\text{Conf}_{t_0,T}^N - \text{Conf}_{t,T}^N) + \overline{P}_{t_0} - \overline{P}_t - \text{PUA}_{[t_0,t)}. \tag{9}$$

Under good event $\mathcal{E}$, $N_{(t_0,t]} - \mathbb{E}\left[N_{(t_0,t]}\right] \geq -\text{Conf}_{t_0,t}^N$. Notice that $\overline{X} \geq \beta_{avg}$ since $\overline{X} = \underline{X} + L_T$ and $L_T \gtrsim \frac{1}{\sqrt{T}}$. Putting this together we have that the left-hand side is lower bounded by $-\beta_{avg}\sqrt{2(t-t_0)|\Theta|\rho_{max}^2\log(2T^2/\delta)} \geq -\beta_{avg}\sqrt{2(T-t_0)|\Theta|\rho_{max}^2\log(2T^2/\delta)}$.

On the other hand, the first term of the right-hand side of (9) is upper bounded by $L_T|\Theta|(\mu_{max} + \rho_{max}) - L_T|\Theta|(t-t_0)$, where we used the fact that $\mathbb{E}[N_t] \geq 1$. Moreover, by definition of the confidence terms and that $\underline{X} \leq \beta_{avg}$, the second term on the right-hand side of (9) is upper bounded by:

$$\beta_{avg}\sqrt{2|\Theta|\rho_{max}^2\log(2T^2/\delta)}\left(\sqrt{T-t_0} - \sqrt{T-t}\right)$$

Lastly, since $\overline{P}_t$ and $\mathrm{PUA}_{[t_0,t)}$ are nonnegative, the last terms on the right hand side are upper bounded by $\overline{P}_{t_0}$. Using the fact that $\overline{P}_{t_0} \leq (1+\sqrt{3\log(2T/\delta)})\sqrt{T-t_0}$ by assumption, we combine these equations and obtain that for any $t \geq t_0$:

$$\begin{aligned}
-\beta_{avg}\sqrt{2(T-t_0)|\Theta|\rho_{max}^2\log(2T^2/\delta)} < L_T|\Theta|(\mu_{max}+\rho_{max}) - L_T|\Theta|(t-t_0) \\
+ \beta_{avg}\sqrt{2|\Theta|\rho_{max}^2\log(2T^2/\delta)}\Big(\sqrt{T-t_0}-\sqrt{T-t}\Big) \\
+ (1+\sqrt{3\log(2T/\delta)})\sqrt{T-t_0}.
\end{aligned}$$

Letting $t = T$ and re-arranging, we then have:

$$\begin{aligned}
0 < L_T|\Theta|(\mu_{max}+\rho_{max}) - L_T|\Theta|(T-t_0) \\
+ \Big(1+\sqrt{3\log(2T/\delta)}+2\beta_{avg}\sqrt{2|\Theta|\rho_{max}^2\log(2T^2/\delta)}\Big)\Big(\sqrt{T-t_0}\Big).
\end{aligned}$$

Let $\xi = 1+\sqrt{3\log(2T/\delta)}+2\beta_{avg}\sqrt{2|\Theta|\rho_{max}^2\log(2T^2/\delta)}$. We will show there exists $c \in \widetilde{\Theta}(1)$ such that, for all $t_0 \leq T - 2cL_T^{-2}$:

$$0 \geq L_T|\Theta|(\rho_{\max}+\mu_{\max}) - L_T|\Theta|(T-t_0) + \xi\sqrt{T-t_0},$$

a contradiction. Denoting $a_1 = |\Theta|(\rho_{\max}+\mu_{\max})$, $a_2 = |\Theta|$, and $a_3 = \xi$ this reduces to showing:

$$0 \geq a_1 L_T - a_2 L_T(T-t_0) + a_3\sqrt{T-t_0}.$$

This is a quadratic inequality in terms of $x = \sqrt{T-t_0}$ with a zero at:

$$\begin{aligned}
x &= \frac{-a_3 - \sqrt{a_3^2+4a_2a_1L_T^2}}{-2a_2L_T} \\
&= \frac{a_3}{2a_2L_T} + \frac{1}{2a_2L_T}\sqrt{a_3^2+4a_2a_1L_T^2} \\
&< \frac{a_3}{2a_2L_T} + \frac{1}{2a_2L_T}\sqrt{a_3^2+4a_2a_1},
\end{aligned}$$

where the final inequality uses the fact that $L_T = o(1)$. Thus, we have shown the existence of $\tilde{c} \in \widetilde{\Theta}(1)$ such that the quadratic is non-positive for any value of

$$x > \tilde{c}L_T^{-1}. \tag{10}$$

Taking this and letting $x = \sqrt{T-t_0}$ shows that the quadratic is negative so long as $T - t_0 \geq \tilde{c}^2 L_T^{-2} \iff t_0 \leq T - \tilde{c}^2 L_T^{-2}$. Relabeling, we have shown that this inequality fails to hold for all $t_0 \leq T - \tilde{c}L_T^{-2}$, a contradiction. □

Proof. Under good event $\mathcal{E}$, we have:

$$P_{<t} \leq \mathbb{E}[P_{<t}] + \mathrm{Conf}_t^P \leq \frac{\mathbb{E}[N_{<t}] - \mathrm{Conf}_{0,t-1}^N}{\mathbb{E}[N] + \mathrm{Conf}_{0,T}^N} B \leq \frac{N_{<t}}{N} B,$$

where the second inequality follows from Assumption (1) of the Theorem, and the third inequality follows again from the fact that, under good event $\mathcal{E}$, $\mathbb{E}[N_{<t}] - \mathrm{Conf}_{0,t-1}^N \leq N_{<t}$, and $N \leq \mathbb{E}[N] + \mathrm{Conf}_{0,T}^N$. Then, by Proposition 3.3, X^{opt} is optimal given $\mathcal{E}$. □

D Section 5 Proofs

D.1 Structure of Offline EG Solutions: Proofs

Proof. Consider any feasible solution $X_{t,\theta}$ to the EG program, and let $\widetilde{X}_{t,\theta} = \frac{\sum_{t'} N_{t',\theta} X_{t',\theta}}{N_\theta}$ for all t. That $\{\widetilde{X}_{t,\theta}\}$ is feasible follows from linearity. We now show that the EG objective under $\{\widetilde{X}_{t,\theta}\}$ is no worse than that under $X_{t,\theta}$.

$$\begin{aligned}\sum_t \sum_\theta N_{t,\theta} \log u(\widetilde{X}_{t,\theta}, \theta) &= \sum_\theta N_\theta \log u\left(\frac{\sum_{t'} N_{t',\theta} X_{t',\theta}}{N_\theta}, \theta\right) \\ &\geq \sum_\theta N_\theta \sum_{t'} \frac{N_{t',\theta}}{N_\theta} \log u(X_{t',\theta}, \theta) \\ &= \sum_t \sum_\theta N_{t,\theta} \log u(X_{t,\theta}, \theta),\end{aligned}$$

where the inequality follows from concavity of the log function and u. □

Proof. For concave, one-positively homothetic, and non-decreasing utilities, any solution to the EG program corresponds to a *competitive equilibrium* in the corresponding Fisher market (cf. Chapter 6 of Nisan et al. (2007)). Formally, let X^* denote an optimal allocation for the EG program, and $p^* = (p^*_k)_{k\in[K]}$ the corresponding optimal dual variables. (X^*, p^*) are said to form a competitive equilibrium if, the following two conditions hold:

1. $\sum_\theta N_\theta X^*_{\theta,k} = B_k$, for all $k \in [K]$, and
2. For all $\theta \in \Theta$, X^*_θ is an optimal solution to:

$$\begin{aligned}\max_{X_\theta \in \mathbb{R}^K_+} \quad & u(X_\theta, \theta) \\ \text{s.t.} \quad & \sum_k X_{\theta k} p^*_k \leq N_\theta,\end{aligned}$$

Moreover, by Theorem 2.2 in Varian (1974), if (X^*, p^*) form a competitive equilibrium, then under this class of functions X^* is a Pareto-efficient and envy-free allocation.

We conclude by showing that X^* is *proportional*. Let $\widetilde{X}$ be such that $\widetilde{X}_\theta = B/N$ for all $\theta \in \Theta$. Clearly, $\widetilde{X}$ is feasible to the EG program. By concavity of u, we have:

$$\begin{aligned}& u(\widetilde{X}_\theta, \theta) \leq u(X^*_\theta, \theta) + \nabla u(X^*_\theta, \theta)^T (\widetilde{X}_{t,\theta} - X^*_\theta) \\ \implies & u(X^*_\theta, \theta) \geq u(\widetilde{X}_\theta, \theta) - \nabla u(X^*_\theta, \theta)^T (\widetilde{X}_\theta - X^*_\theta).\end{aligned}$$

Using the fact that X^*_θ is the utility-maximizing allocation for type $\theta \in \Theta$ and concavity of u, we have: $\nabla u(X^*_\theta, \theta)^T (\widetilde{X}_\theta - X^*_\theta) \leq 0$. Substituting in $\widetilde{X}_\theta = B/N$, we obtain the result. □

Proof. Let $(p_k)_{k\in[K]}$ and $(\alpha_{\theta,k})_{\theta\in\Theta,k\in[K]}$ respectively denote optimal dual solutions for the EG program for an arbitrary vector of arrivals. The KKT conditions for the EG program are given by:

1. *Primal Feasibility*: $\sum_{\theta\in\Theta} N_\theta X_{\theta,k} \leq B_k$ for all k, $X_{\theta,k} \geq 0$ for all θ, k
2. *Dual Feasibility*: $p_k \geq 0$ for all k, $\alpha_{\theta,k} \geq 0$ for all θ, k

3. *Complementary Slackness*:

$$\begin{cases} p_k > 0 \implies \sum_{\theta\in\Theta} N_\theta X_{\theta,k} = B_k \\ \alpha_{\theta,k} > 0 \implies X_{\theta,k} \;=\; 0 \end{cases}$$

4. *Gradient Condition*:

$$\frac{-N_\theta \frac{\partial}{\partial X_{\theta,k}} u(X_\theta,\theta)}{u(X_\theta,\theta)} + N_\theta p_k - \alpha_{\theta,k} = 0 \quad \forall\, \theta \in \Theta, k \in [K].$$

Using the fact that $\alpha_{\theta,k} \geq 0$, with equality if $X_{\theta,k} > 0$ we find that:

$$p_k \geq \frac{\frac{\partial}{\partial X_{\theta,k}} u(X_\theta,\theta)}{u(X_\theta,\theta)} \quad \forall\, \theta \in \Theta, k \in [K]. \tag{11}$$

with equality whenever $X_{\theta,k} > 0$.

We now use these conditions to prove the scaling and monotonicity properties. For ease of notation, we let X^* and (p^*, α^*) denote optimal primal and dual solutions corresponding to arrivals $(N_\theta)_{\theta\in\Theta}$.

Scaling: Suppose $\widetilde{N}_\theta = (1+\zeta)N_\theta$ for every $\theta \in \Theta$.

For all $k \in [K]$, $\theta \in \Theta$, define $\tilde{p}_k = (1+\zeta)p_k^*$, $\widetilde{X}_{\theta,k} = X^*_{\theta,k}/(1+\zeta)$, and $\tilde{\alpha}_{\theta,k} = (1+\zeta)^2\alpha^*_{\theta,k}$. Since our constraint set is linear, it suffices to check that $\widetilde{X}$, $\tilde{p}$ satisfy the KKT conditions.

Dual feasibility follows from the fact that $\zeta \geq 0$, and primal feasibility follows from:

$$\sum_\theta \widetilde{N}_\theta \widetilde{X}_{\theta,k} = \sum_\theta (1+\zeta)N_\theta \frac{X^*_{\theta,k}}{1+\zeta} \leq B_k$$

by feasibility of $X^*_{\theta,k}$.

Complementary slackness holds since the resource utilizations are equal under both arrival vectors, and $p_k^* > 0 \iff \tilde{p}_k > 0$. A similar argument holds for $\tilde{\alpha}_{\theta,k}$.

Finally, we verify the gradient condition. Since u is one-positively homogeneous, it follows that its partial derivatives are 0-positively homogeneous, and as a result $\left.\frac{\partial u}{\partial X_{\theta,k}}\right|_{X_\theta=\widetilde{X}_\theta} = \left.\frac{\partial u}{\partial X_{\theta,k}}\right|_{X_\theta=X^*_\theta}$.

Thus, we have that:

$$\begin{aligned}
\frac{-\widetilde{N}_\theta \frac{\partial}{\partial X_{\theta,k}} u(X_\theta,\theta)\big|_{X_\theta=\widetilde{X}_\theta}}{u(\widetilde{X}_\theta,\theta)} + \widetilde{N}_\theta \tilde{p}_k - \tilde{\alpha}_{\theta,k} &= \frac{-(1+\zeta)N_\theta \frac{\partial}{\partial X_{\theta,k}} u(X_\theta,\theta)\big|_{X_\theta=X^*_\theta}}{\frac{1}{1+\zeta}u(X_\theta,\theta)} + (1+\zeta)^2 N_\theta p_k \\
&\quad - (1+\zeta)^2\alpha_{\theta,k} \\
&= (1+\zeta)^2\left(\frac{-N_\theta \frac{\partial}{\partial X_{\theta,k}} u(X_\theta,\theta)\big|_{X_\theta=X^*_\theta}}{u(X_\theta,\theta)} + N_\theta p_k - \alpha_{\theta,k}\right) \\
&= 0
\end{aligned}$$

by optimality of $X^*_{\theta,k}, p^*_k, \alpha^*_{\theta,k}$.

Plugging optimality of $\widetilde{X}_\theta = X^*_\theta/(1+\zeta)$, we obtain:

$$u(X^*_\theta,\theta) - u(\widetilde{X}_\theta,\theta) = u(X^*_\theta,\theta) - u(X^*_\theta/(1+\zeta),\theta) = \left(1 - \frac{1}{1+\zeta}\right)u(X^*_\theta,\theta),$$

where the final equality follows from the fact that the utility functions are one-homogeneous.

Monotonicity: As before, let X^* and $\widetilde{X}$ respectively denote optimal solutions to the EG program under $(N_\theta)_{\theta\in\Theta}$ and $(\widetilde{N}_\theta)_{\theta\in\Theta}$ arrivals, with $\widetilde{N}_\theta \geq N_\theta$ for all θ, and a strict inequality for at least one type θ. Similarly, let $(p_k^*)_{k\in[K]}$ and $(\widetilde{p}_k)_{k\in[K]}$ be the corresponding dual solutions.

Suppose for contradiction that there exists a type $\theta \in \Theta$ such that $u(\widetilde{X}_\theta, \theta) > u(X_\theta^*, \theta)$. Since $u(\cdot, \theta)$ is non-decreasing, it must be that $\widetilde{X}_{\theta,k} > X^*_{\theta,k}$ for some $k \in [K]$. By the KKT condition (11), we have:

$$u(\widetilde{X}_\theta, \theta) = \frac{\frac{\partial}{\partial X_{\theta,k}} u(X_\theta, \theta)\big|_{X_\theta = \widetilde{X}_\theta}}{\widetilde{p}_k} > u(X_\theta^*, \theta) \geq \frac{\frac{\partial}{\partial X_{\theta,k}} u(X_\theta, \theta)\big|_{X_\theta = X_\theta^*}}{p_k^*}.$$

Moreover, since u is concave and $\widetilde{X}_{\theta,k} > X^*_{\theta,k}$, $\frac{\partial}{\partial X_{\theta,k}} u(X_\theta, \theta)\big|_{X_\theta = \widetilde{X}_\theta} \leq \frac{\partial}{\partial X_{\theta,k}} u(X_\theta, \theta)\big|_{X_\theta = X_\theta^*}$. Thus, it must be that $\widetilde{p}_k < p_k^*$.

Now, note that, since $\widetilde{X}_{\theta,k} > X^*_{\theta,k}$, it must be that $\widetilde{X}_{\theta',k} < X^*_{\theta',k}$ for some $\theta' \in \Theta$ since X^* clears the resources B_k. Again, using concavity of u, this implies that $\frac{\partial}{\partial X_{\theta,k}} u(X_{\theta'}, \theta')\big|_{X_{\theta'} = \widetilde{X}_{\theta'}} \geq \frac{\partial}{\partial X_{\theta,k}} u(X_{\theta'}, \theta')\big|_{X_{\theta'} = X^*_{\theta'}}$. Putting this together with the fact that $\widetilde{p}_k < p_k^*$, we obtain:

$$\frac{\frac{\partial}{\partial X_{\theta',k}} u(X_{\theta'}, \theta')\big|_{X_{\theta'} = \widetilde{X}_{\theta'}}}{\widetilde{p}_k} > \frac{\frac{\partial}{\partial X_{\theta',k}} u(X_{\theta'}, \theta')\big|_{X_{\theta'} = X^*_{\theta'}}}{p_k^*},$$

and by the KKT condition (11), $u(\widetilde{X}_{\theta'}, \theta') > u(X^*_{\theta'}, \theta')$.

Note however that, since $(\widetilde{N}_\theta)_{\theta\in\Theta} \geq (N_\theta)_{\theta\in\Theta}$, $\widetilde{X}$ is feasible to the EG program under $(N_\theta)_{\theta\in\Theta}$. Thus, the objective for this latter program under $\widetilde{X}$ *strictly* improves upon the objective under X^*, which contradicts optimality of X^*. □

D.2 Theorem 5.4

In the remainder of the section, for ease of notation we let $\textsc{Conf}_\theta = \textsc{Conf}_{0,T,\theta}$. Moreover, we define constants $C_1 = \max_{\theta\in\Theta} u(B, \theta), C_2 = B_{\min}/\mathbb{E}[N], C_3 = \|B\|_\infty$, where $B_{\min} = \min_{k\in[K]} B_k$.

The following proposition will be of use throughout the proof. We defer its proof to Appendix D.2.2.

Proposition D.1. *Under good event $\mathcal{E}$, $C_1 \geq \max_{\theta\in\Theta} u(\overline{X}, \theta)$, $C_2 \leq \min_{\theta\in\Theta} \|\overline{X}_\theta\|_\infty$, $C_3 \geq \max_{\theta\in\Theta} \|\overline{X}_\theta\|_\infty$.*

We now state the main building blocks of our result. Lemma D.2 implies that it suffices to restrict our attention to $\mathcal{E}$ for the high-probability statement. We omit the this result, as it is entirely analogous to that of Lemma C.1.

Lemma D.2. *With probability at least $1-\delta$, for every $\theta \in \Theta$ and $t' > t$, $|N_{(t,t'],\theta} - \mathbb{E}[N_{(t,t'],\theta}]| \leq \textsc{Conf}_{t,t',\theta}$.*

Next, Lemmas D.3 and D.4 formalize both that NL-Guardrail only ever allocates the guardrails under $\mathcal{E}$, and that it is adaptively cautious.

Lemma D.3. *Given event $\mathcal{E}$, for every resource k and time $t \in [T]$,*

$$B^{alg}_{t,k} \geq \sum_{\theta\in\Theta} N_{\geq t,\theta} \underline{X}_{\theta,k}.$$

As a result, given $\mathcal{E}$, $X^{alg}_{t,\theta,k} \in \{\overline{X}_{\theta,k}, \underline{X}_{\theta,k}\}$ for all $t \in [T], \theta \in \Theta, k \in [K]$.

Lemma D.4. *Consider event $\mathcal{E}$, and let $t_{0,k}$ be the last round for which $X^{alg}_{t,\theta,k} = \underline{X}_{\theta,k}$, (else, define $t_{0,k} = 0$ if $X^{alg}_{t,\theta,k} = \underline{X}_{\theta,k}$ for all $t \in [T]$) for each resource $k \in [K]$. Then, $t_{0,k} \geq \max\{0, T - cL_T^{-2}\}$ for all $k \in [K]$, where $c = \widetilde{\Theta}(1)$.*

The following lemma establishes that agents' utilities under two different allocations differ by at most L_T, and that the allocations themselves are within a factor of L_T of each other. Moreover, under a lower bound on L_T, not only is the true vector of arrivals sandwiched by the arrival confidence bounds, but the hindsight optimal utility is sandwiched by the guardrail utilities.

Lemma D.5. *The following holds for our guardrail allocations $\underline{X}_\theta$ and $\overline{X}_\theta$:*

1. $u(x((\underline{n}_\theta)_{\theta\in\theta})_\theta, \theta) - u(x((\overline{n}_\theta)_{\theta\in\theta})_\theta, \theta) \leq L_T$
2. $\|x((\overline{n}_\theta)_{\theta\in\Theta}) - x((\underline{n}_\theta)_{\theta\in\theta})\|_\infty \geq \frac{C_2}{C_1} L_T$
3. $\|x((\overline{n}_\theta)_{\theta\in\Theta}) - x((\underline{n}_\theta)_{\theta\in\theta})\|_\infty \leq \frac{C_3}{C_1} L_T$.

If in addition $L_T \geq 2C_1 \max_\theta \frac{\textsc{Conf}_\theta}{\mathbb{E}[N_\theta]}$, under event $\mathcal{E}$ we have:

4. $\underline{n}_\theta \leq N_\theta \leq \overline{n}_\theta$ *for all* $\theta \in \Theta$
5. $u(x((\overline{n}_\theta)_{\theta\in\theta})_\theta, \theta) \leq u(X^{opt}_\theta, \theta) \leq u(x((\underline{n}_\theta)_{\theta\in\theta})_\theta, \theta)$.

Given these building blocks of our main result, we now prove the theorem.

Proof. **Efficiency**: Consider first $\Delta_{\mathit{efficiency}}$. We have:

$$
\begin{aligned}
\Delta_{\mathit{efficiency}} &= \sum_k \left(B_k - \sum_{t\in[T]} \sum_\theta X^{alg}_{t,\theta,k} N_{t,\theta} \right) = \sum_k \left(B^{alg}_{t_{0,k}} - \sum_{t\geq t_{0,k}} \sum_{\theta\in\Theta} N_{t,\theta} X^{alg}_{t,\theta,k} \right) \\
&\overset{(a)}{=} \sum_k B^{alg}_{t_{0,k},k} - \sum_\theta \left(\underline{X}_{\theta,k} N_{t_{0,k},\theta} + \sum_{t>t_{0,k}} \overline{X}_{\theta,k} N_{t,\theta} \right) \\
&\overset{(b)}{<} \sum_k \sum_\theta N_{t_{0,k},\theta} \overline{X}_{\theta,k} + \sum_\theta \underline{X}_{\theta,k} \big(\mathbb{E}\big[N_{>t_{0,k},\theta}\big] + \textsc{Conf}_{t_{0,k},T,\theta} \big) \\
&\qquad - \sum_\theta \left(\underline{X}_{\theta,k} N_{t_{0,k},\theta} + \sum_{t>t_{0,k}} \overline{X}_{\theta,k} N_{t,\theta} \right) \\
&= \sum_k \sum_\theta \underline{X}_{\theta,k} \big(\mathbb{E}\big[N_{>t_{0,k},\theta}\big] + \textsc{Conf}_{t_{0,k},T,\theta} - N_{>t_{0,k},\theta} \big) - (\overline{X}_{\theta,k} - \underline{X}_{\theta,k})(N_{>t_{0,k},\theta} - N_{t_{0,k},\theta}) \\
&\overset{(c)}{\leq} 2 \sum_k \sum_\theta \underline{X}_{\theta,k} \textsc{Conf}_{t_{0,k},T,\theta} + N_{t_{0,k},\theta} (\overline{X}_{\theta,k} - \underline{X}_{\theta,k})
\end{aligned}
$$

where (a) follows from the fact that, by the definition of $t_{0,k}$, $X^{alg}_{t_{0,k},\theta,k} = \underline{X}_{\theta,k}$, and $X^{alg}_{t,\theta,k} = \overline{X}_{\theta,k}$ for all $t > t_{0,k}$; (b) follows from the condition in the algorithm for allocating the lower allocation at time $t_{0,k}$, which upper bounds $B^{alg}_{t_{0,k},k}$; and (c) follows from the fact that, under $\mathcal{E}$, $\mathbb{E}\big[N_{>t_{0,k},\theta}\big] - N_{>t_{0,k},\theta} \leq \textsc{Conf}_{t_{0,k},T,\theta}$.

Plugging in the definition of $\textsc{Conf}_{t_{0,k},T,\theta}$ and the fact that $(\overline{X}_{\theta,k} - \underline{X}_{\theta,k}) \leq \frac{C_3}{C_1} L_T$ by Lemma D.5, we have:

$$\Delta_{\mathit{efficiency}} \leq 2C_3 \sum_k \sum_\theta \sqrt{2\rho_{max}^2 (T - t_{0,k}) \log(2T^2|\Theta|/\delta)} + (\mu_{max} + \rho_{max}) \frac{C_3}{C_1} L_T.$$

Taking this and plugging in the lower bound on $t_{0,k}$ from Lemma D.4, we get that:

$$\Delta_{\mathit{efficiency}} \leq 2C_3 K|\Theta| \sqrt{2\rho_{max}^2 \log(2T^2|\Theta|/\delta)} \min\{\sqrt{T}, \sqrt{2c}/L_T\} + K|\Theta|(\mu_{max} + \rho_{max}) \frac{C_3}{C_1} L_T.$$

Note that, for $L_T = o(1)$, the second term is dominated by the first, which gives us the desired bound on $\Delta_{\mathit{efficiency}}$.

Hindsight Envy: Fix $t, t' \in [T]$, and $\theta, \theta' \in \Theta$. By Lemma D.3, $X_{t,\theta,k}^{alg} \in \{\overline{X}_{\theta,k}, \underline{X}_{\theta,k}\}$ for all $t \in [T], \theta \in \Theta, k \in [K]$. Using the fact that u is non-decreasing, we have:

$$\begin{aligned} u(X_{t',\theta'}^{alg}, \theta) - u(X_{t,\theta}^{alg}, \theta) &\leq u(\overline{X}_{\theta'}, \theta) - u(\underline{X}_\theta, \theta) \\ &= u(\overline{X}_{\theta'}, \theta) - u(\overline{X}_\theta, \theta) + u(\overline{X}_\theta, \theta) - u(\underline{X}_\theta, \theta) \\ &\overset{(a)}{\leq} u(\overline{X}_\theta, \theta) - u(\underline{X}_\theta, \theta) \\ &\overset{(b)}{\leq} L_T, \end{aligned}$$

where (a) follows from the fact that $\overline{X}$ is envy-free by construction, and (b) follows from the bound in Lemma D.5. Taking the max over t, t', θ, θ' gives the result.

Counterfactual Envy: We next obtain the bound on Δ_{EF}. Consider first the setting where $L_T \geq 2C_1 \sqrt{\frac{2\rho_{max}^2 \log(2T^2|\Theta|/\delta)}{T}}$, such that Properties 4 and 5 of Lemma D.5 are satisfied. Again, using Lemma D.3 along with the fact that $u(\underline{X}, \theta) \leq u(X^{opt}, \theta) \leq u(\overline{X}, \theta)$, by Lemma D.5, we have that, given $\mathcal{E}$:

$$|u(X_{t,\theta}^{alg}, \theta) - u(X_\theta^{opt}, \theta)| \leq |u(\overline{X}_\theta, \theta) - u(\underline{X}_\theta, \theta)| \leq L_T.$$

Consider now the case where $L_T < 2C_1 \sqrt{\frac{2\rho_{max}^2 \log(2T^2|\Theta|/\delta)}{T}}$ then, one of the following two chain of inequalities holds for all $\theta \in \Theta$: (1) $u(\underline{X}_\theta, \theta) \leq u(X_\theta^{opt}, \theta) \leq u(\overline{X}_\theta, \theta)$, or (2) $u(\underline{X}_\theta, \theta) \leq u(\overline{X}_\theta, \theta) \leq u(X_\theta^{opt}, \theta)$, since u is non-decreasing. In the first case, we have $|u(X_{t,\theta}^{alg}, \theta) - u(X_\theta^{opt}, \theta)| \leq L_T$ as above.

Else, let $\widetilde{X}_\theta$ be the upper guardrail solution corresponding to $\widetilde{L}_T = 2C_1 \sqrt{\frac{2\rho_{max}^2 \log(2T^2|\Theta|/\delta)}{T}}$. Via the same reasoning as above, we have:

$$|u(X_{t,\theta}^{alg}, \theta) - u(X_\theta^{opt}, \theta)| \leq |u(\widetilde{X}_\theta, \theta) - u(\underline{X}_\theta, \theta)| \leq \widetilde{L}_T = 2C_1 \sqrt{\frac{2\rho_{max}^2 \log(2T^2|\Theta|/\delta)}{T}},$$

where the second inequality again follows from Lemma D.5.

Hindsight Proportionality: We conclude by showing the bound on Δ_{prop}. We have:

$$u(B/N, \theta) - u(X_{t,\theta}^{alg}, \theta) = u(B/N, \theta) - u(X_\theta^{opt}, \theta) + u(X_\theta^{opt}, \theta) - u(X_{t,\theta}^{alg}, \theta) \lesssim \max\{1/\sqrt{T}, L_T\},$$

where the final inequality follows from the fact that X^{opt} satisfies proportionality, by Proposition 5.2, and thus $u(B/N, \theta) - u(X_\theta^{opt}, \theta) \leq 0$, and $u(X_\theta^{opt}, \theta) - u(X_{t,\theta}^{alg}, \theta) \leq \Delta_{EF} \lesssim \max\{1/\sqrt{T}, L_T\}$, which we just proved above. Taking the max over t and θ gives the desired bound on Δ_{prop}. □

D.2.1 Proofs of main building blocks

Proof. We show the first statement by induction on t. Given the first statement, the second follows by construction.

Base Case: $t = 1$. We first note that $B_1^{alg} = B$. Moreover, since $\underline{X} = x(\overline{n}_\theta)$, we have that, for all $k \in [K]$:

$$\sum_\theta \overline{n}_\theta \underline{X}_{\theta,k} \le B_{1,k}^{alg}.$$

Plugging in the definition of $\overline{n}_\theta = \mathbb{E}[N_\theta](1 + \max_\theta \frac{\textsc{Conf}_\theta}{\mathbb{E}[N_\theta]})$, we obtain:

$$\begin{aligned} B_{1,k}^{alg} &\ge \sum_\theta \left(\mathbb{E}[N_\theta] \left(1 + \max_\theta \frac{\textsc{Conf}_\theta}{\mathbb{E}[N_\theta]} \right) \right) \underline{X}_{\theta,k} \\ &\ge \sum_\theta (\mathbb{E}[N_\theta] + \textsc{Conf}_\theta) \underline{X}_{\theta,k} \\ &\ge \sum_\theta N_\theta \underline{X}_{\theta,k}, \end{aligned}$$

where the final inequality follows from the definition of event $\mathcal{E}$.

Step Case: $t - 1 \to t$. We split the analysis into two cases, based on the allocation in round $t - 1$.

If $X_{t-1,\theta,k}^{alg} = \underline{X}_{\theta,k}$, then

$$B_{t,k}^{alg} = B_{t-1,k}^{alg} - \sum_{\theta\in\Theta} N_{t-1,\theta} \underline{X}_{\theta,k} \ge \sum_{\theta\in\Theta} N_{\ge t,\theta} \underline{X}_{\theta,k},$$

where the last inequality follows from the induction hypothesis.

If $X_{t-1,\theta,k}^{alg} = \overline{X}_{\theta,k}$, then

$$B_{t,k}^{alg} = B_{t-1,k}^{alg} - \sum_{\theta\in\Theta} N_{t-1,\theta} X_{t-1,\theta,k}^{alg} \overset{(a)}{\ge} \sum_{\theta\in\Theta} \underline{X}_{\theta,k} (\mathbb{E}[N_{\ge t,\theta}] + \textsc{Conf}_{t-1,T,\theta}) \overset{(b)}{\ge} \sum_{\theta\in\Theta} N_{\ge t,\theta} \underline{X}_{\theta,k}.$$

where (a) holds by the condition for allocating $\overline{X}_{\theta,k}$, and (b) holds under event $\mathcal{E}$. □

Proof. By definition of $t_{0,k}$, it must be that $X_{t_{0,k},\theta,k}^{alg} = \underline{X}_{\theta,k}$, and $X_{t,\theta,k}^{alg} = \overline{X}_{\theta,k}$ for all $t > t_{0,k}$. Thus, for all $t > t_{0,k}$:

$$\begin{aligned} \sum_\theta \underline{X}_{\theta,k} (\mathbb{E}[N_{>t,\theta}] + \textsc{Conf}_{t,T,\theta}) &\overset{(a)}{\le} B_{t,k}^{alg} - \sum_\theta N_{t,\theta} \overline{X}_{\theta,k} \overset{(b)}{=} B_{t_{0,k}}^{alg} - \sum_\theta \underline{X}_{\theta,k} N_{t_{0,k},\theta} - \sum_\theta \overline{X}_{\theta,k} \sum_{t'=t_{0,k}+1}^{t} N_{t',\theta} \\ &\overset{(c)}{<} \sum_\theta \underline{X}_{\theta,k} \big(\mathbb{E}\big[N_{>t_{0,k},\theta}\big] + \textsc{Conf}_{t_{0,k},T,\theta} \big) + \sum_\theta (\overline{X}_{\theta,k} - \underline{X}_{\theta,k}) N_{t_{0,k},\theta} - \sum_\theta \overline{X}_{\theta,k} \sum_{t'=t_{0,k}+1}^{t} N_{t',\theta} \end{aligned}$$

where (a) follows from the condition in the algorithm for $X_{t,\theta,k}^{alg} = \overline{X}_{\theta,k}$, (b) follows from the definition of $t_{0,k}$ and the choice of $t > t_{0,k}$, and (c) follows from the condition in the algorithm for $X_{t_{0,k},\theta,k}^{alg} = \underline{X}_{\theta,k}$.

Let $N_{(t_{0,k},t],\theta} = \sum_\theta \sum_{t'=t_{0,k}+1}^{t} N_{t',\theta}$. Re-arranging the above inequality, we obtain:

$$\sum_\theta \overline{X}_{\theta,k} N_{(t_{0,k},t],\theta} < \sum_\theta (\overline{X}_{\theta,k} - \underline{X}_{\theta,k}) N_{t_{0,k},\theta} + \sum_\theta \underline{X}_{\theta,k} \mathbb{E}\Big[N_{(t_{0,k},t],\theta}\Big] + \sum_\theta \underline{X}_{\theta,k}\big(\textsc{Conf}_{t_{0,k},T,\theta} - \textsc{Conf}_{t,T,\theta}\big)$$

$$\iff \sum_\theta \overline{X}_{\theta,k}(N_{(t_{0,k},t],\theta} - \mathbb{E}\Big[N_{(t_{0,k},t],\theta}\Big]) < \sum_\theta (\overline{X}_{\theta,k} - \underline{X}_{\theta,k})\Big(N_{t_{0,k},\theta} - \mathbb{E}\Big[N_{(t_{0,k},t],\theta}\Big]\Big) + \sum_\theta \underline{X}_{\theta,k}\big(\textsc{Conf}_{t_{0,k},T,\theta} - \textsc{Conf}_{t,T,\theta}\big) \tag{12}$$

Note that $\frac{C_2}{C_1} L_T \le \overline{X}_{\theta,k} - \underline{X}_{\theta,k} \le \frac{C_3}{C_1} L_T$ by Lemma D.5. Moreover, $|N_{t_{0,k},\theta} - \mathbb{E}[N_{t_{0,k},\theta}]| \le \rho_{\max} \implies N_{t_{0,k},\theta} \le \mu_{\max} + \rho_{\max}$. We moreover use the fact that $\mathbb{E}[N_{t,\theta}] \ge 1$ for all t, θ to obtain the following upper bound on the first term of (12):

$$\sum_\theta (\overline{X}_{\theta,k} - \underline{X}_{\theta,k})\Big(N_{t_{0,k},\theta} - \mathbb{E}\Big[N_{(t_{0,k},t],\theta}\Big]\Big) \le \frac{C_3}{C_1} L_T |\Theta| (\rho_{\max} + \mu_{\max}) - \frac{C_2}{C_1} L_T |\Theta| (t - t_{0,k}).$$

Putting this together with the definition of the confidence terms and the absolute upper bound C_3 on $||\underline{X}_{\theta,k}||_\infty$, we obtain:

$$\sum_\theta \overline{X}_{\theta,k}\Big(N_{(t_{0,k},t],\theta} - \mathbb{E}\Big[N_{(t_{0,k},t],\theta}\Big]\Big) < \frac{C_3}{C_1} L_T |\Theta| (\rho_{\max} + \mu_{\max}) - \frac{C_2}{C_1} L_T |\Theta| (t - t_{0,k}) + C_3 |\Theta| \sqrt{2\rho_{max}^2 \log(2T^2|\Theta|/\delta)}\Big(\sqrt{T - t_{0,k}} - \sqrt{T - t}\Big).$$

Moreover, given event $\mathcal{E}$, we have the following lower bound on the left-hand side of (12):

$$\begin{aligned}\sum_\theta \overline{X}_{\theta,k}(N_{(t_{0,k},t],\theta} - \mathbb{E}\Big[N_{(t_{0,k},t],\theta}\Big]) &\ge -C_3|\Theta|\sqrt{2\rho_{max}^2 \log(2T^2\Theta|/\delta)(t - t_{0,k})} \\ &\ge -C_3|\Theta|\sqrt{2\rho_{max}^2 \log(2T^2\Theta|/\delta)(T - t_{0,k})},\end{aligned}$$

where the final inequality is simply used to simplify algebraic manipulations later on. Combining these equations we obtain that for any $t \ge t_{0,k}$:

$$-C_3|\Theta|\sqrt{2\rho_{max}^2 \log(2T^2\Theta|/\delta)(T - t_{0,k})} < \frac{C_3}{C_1} L_T |\Theta| (\rho_{\max} + \mu_{\max}) - \frac{C_2}{C_1} L_T |\Theta| (t - t_{0,k}) + C_3 |\Theta| \sqrt{2\rho_{max}^2 \log(2T^2|\Theta|/\delta)}\Big(\sqrt{T - t_{0,k}} - \sqrt{T - t}\Big).$$

We omit the remainder of the proof follows entirely analogous lines as that of Lemma 4.4. Namely, we prove the existence of $c \in \widetilde{\Theta}(1)$ such that for all $t_{0,k} \le T - cL_T^{-2}$, this inequality fails to hold. □

Proof. For ease of notation, let $\gamma = \max_\theta \textsc{Conf}_\theta / \mathbb{E}[N_\theta]$. Then $\overline{n}_\theta = \mathbb{E}[N_\theta](1+\gamma)$, $\underline{n}_\theta = \mathbb{E}[N_\theta](1-c)$, and so $\overline{n}_\theta = \frac{1+\gamma}{1-c}\underline{n}_\theta$. The result follows from the following properties of solutions to the offline EG program, established in Proposition 5.3.

We first show Parts 1-3 of the claim. By Part 1 of Proposition 5.3, we have that:

$$u(x((\underline{n}_\theta)_{\theta\in\theta})_\theta, \theta) - u(x((\overline{n}_\theta)_{\theta\in\theta})_\theta, \theta) = \Big(1 - \frac{1-c}{1+\gamma}\Big) u(x((\underline{n}_\theta)_{\theta\in\theta})_\theta, \theta) \le \frac{c+\gamma}{1+\gamma} C_1 = L_T.$$

Using Part 1 of Proposition 5.3 once again we can lower bound the difference in allocations by

$$\|x((\overline{n}_\theta)_{\theta\in\Theta}) - x((\underline{n}_\theta)_{\theta\in\theta})\|_\infty = \left(1 - \frac{1-c}{1+\gamma}\right)\|x((\underline{n}_\theta)_{\theta\in\theta})\|_\infty \geq \frac{c+\gamma}{1+\gamma}C_2 = \frac{C_2}{C_1}L_T.$$

The argument for the upper bound is identical (where instead we use $\|x((\underline{n}_\theta)_{\theta\in\theta})\|_\infty \leq C_3$).

We now show Parts 4 and 5 of the claim. We have:

$$N_\theta = \mathbb{E}[N_\theta] + (N_\theta - \mathbb{E}[N_\theta]) = \mathbb{E}[N_\theta]\left(1 + \frac{N_\theta - \mathbb{E}[N_\theta]}{\mathbb{E}[N_\theta]}\right) \tag{13}$$

Under event $\mathcal{E}$, $\textsc{Conf}_\theta$ upper bounds $|N_\theta - \mathbb{E}[N_\theta]|$. Plugging this into Eq. (13), we obtain that $N_\theta \leq \mathbb{E}[N_\theta]\left(1 + \max_\theta \frac{\textsc{Conf}_\theta}{\mathbb{E}[N_\theta]}\right) = \overline{n}_\theta$ given $\mathcal{E}$.

Via symmetric reasoning, $N_\theta \geq \mathbb{E}[N_\theta]\left(1 - \max_\theta \frac{\textsc{Conf}_\theta}{\mathbb{E}[N_\theta]}\right)$ given $\mathcal{E}$. Simple algebraic manipulations and the assumption that $L_T \geq 2C_1 \max_\theta \frac{\textsc{Conf}_\theta}{\mathbb{E}[N_\theta]}$ show that $\mathbb{E}[N_\theta]\left(1 - \max_\theta \frac{\textsc{Conf}_\theta}{\mathbb{E}[N_\theta]}\right) \geq \underline{n}_\theta$, and thus $N_\theta \geq \underline{n}_\theta$.

Finally, the fact that $u(x((\overline{n}_\theta)_{\theta\in\Theta})_\theta, \theta) \leq u(X_\theta^{opt}, \theta) \leq u(x((\underline{n}_\theta)_{\theta\in\theta})_\theta, \theta)$ given $\mathcal{E}$ follows from putting Part 4 of the above claim and the monotonicity property in Proposition 5.3 together. □

D.2.2 Proofs of Auxiliary Results

Proof. For C_1, note that $\overline{X} \leq B$ by the budget constraint, so $\max_{\theta\in\Theta} u(\overline{X}_\theta, \theta) \leq \max_{\theta\in\Theta} u(B, \theta)$, since u is non-decreasing.

For C_2 we use that $\overline{X}$ satisfies $B_k = \sum_\theta \overline{X}_{\theta,k}\underline{n}_\theta$ for every θ by Pareto efficiency of the optimal solutions to the EG program, and the assumption that the utilities are strictly monotone. Using Cauchy-Schwarz and the fact that $\underline{n}_\theta \leq \mathbb{E}[N_\theta]$ we notice:

$$\begin{aligned} B_k &= \sum_\theta \underline{n}_\theta \overline{X}_{\theta,k} \leq (\max_\theta \overline{X}_{\theta,k}) \sum_\theta \underline{n}_\theta = \|\overline{X}_\theta\|_\infty \mathbb{E}[N] \\ \implies \|\overline{X}_\theta\|_\infty &\geq B_k/\mathbb{E}[N] \geq B_{\min}/\mathbb{E}[N]. \end{aligned}$$

Lastly, for C_3 we note again by the budget constraint that $\max_{\theta\in\Theta}\|\overline{X}_\theta\|_\infty \leq \|B\|_\infty$. □

E Leontief Utilities

In this section, leveraging the observation that ϵ-perturbed Leontief utilities satisfy the requirements for our main results of Section 5, we obtain bounds on $(\Delta_{EF}, \Delta_{\mathit{efficiency}}, \textsc{Envy})$ for Leontief utilities, given by:

$$u(x, \theta) = \min\{x_k/w_{\theta,k}\} \quad \forall\, \theta \in \Theta.$$

Let $u_\epsilon(x, \theta) = u(x, \theta) + \epsilon\langle w'_\theta, x\rangle$ for some $\epsilon > 0$. We refer to $u(x, \theta)$ as the *nominal* utility function and $u_\epsilon(x, \theta)$ as the *perturbed* utility function. Finally, we introduce the following notions of ϵ-perturbed fairness.

Definition E.1 (Perturbed Counterfactual Envy, Hindsight Envy, and Efficiency)**.** *Given individuals with types Θ, sizes $(N_{t,\theta})_{t\in[T],\theta\in\Theta}$, and resource budgets $(B_k)_{k\in[K]}$, for any online allocation $(X_{t,\theta}^{alg})_{t\in[T],\theta\in\Theta} \in \mathbb{R}^k$, we define:*

- Counterfactual Envy*: The counterfactual distance of X^{alg} to envy-freeness as*

$$\Delta_{EF}(\epsilon_1, \epsilon_2) \triangleq \max_{t\in[T],\theta\in\Theta} \|u_{\epsilon_1}(X_{t,\theta}^{alg}, \theta) - u_{\epsilon_1}(X_{t,\theta}^{opt(\epsilon_2)}, \theta)\|_\infty$$

where $X^{opt(\epsilon_2)}$ is the solution to (EG) with true values $(N_{t,\theta})_{t\in[T],\theta\in\Theta}$ and utilities u_{ϵ_2}.

- Hindsight Envy*: The hindsight distance of X^{alg} to envy-freeness as*

$$\textsc{Envy}(\epsilon) \triangleq \max_{t,t'\in[T]^2,\theta,\theta'\in\Theta^2} u_\epsilon(X_{t',\theta'}^{alg}, \theta) - u_\epsilon(X_{t,\theta}^{alg}, \theta).$$

- Efficiency*: The distance to efficiency as*

$$\Delta_{\mathit{efficiency}} \triangleq \sum_{k\in K}\left(B_k - \sum_{t\in[T]}\sum_{\theta\in\Theta} N_{t,\theta} X_{t,\theta,k}^{alg}\right)$$

Given these definitions, obtaining guarantees with respect to the *nominal* Leontief utilities is equivalent to obtaining ϵ-perturbed fairness guarantees on $\Delta_{EF}(0,\epsilon), \textsc{Envy}(0)$, and $\Delta_{\mathit{efficiency}}$, for some $\epsilon > 0$. Note that we use $X^{opt(\epsilon)}$ as the benchmark allocation in this setting, as in general solutions to EG under nominal utilities may not be fair without strict monotonicity.

Given this setup, the following theorem shows that we are able to obtain the desired bounds on the nominal utilities by running NL-Guardrail on the perturbed utilities, for an appropriately chosen value of ϵ.

Theorem E.2. *Fix $L_T = o(1)$, and let* NL-Guardrail *be initialized with utility functions u_ϵ where $\epsilon = L_T$, and*

$$\begin{aligned}
\textsc{Conf}_{t,t',\theta} &= \sqrt{2(t'-t)\rho_{max}^2 \log(2T^2|\Theta|/\delta)} \quad \forall t' > t \in [T], \theta \in \Theta \\
\overline{n}_\theta &= \mathbb{E}[N_\theta]\left(1 + \max_\theta \frac{\textsc{Conf}_{0,T,\theta}}{\mathbb{E}[N_\theta]}\right) \quad \forall \theta \in \Theta \\
\underline{n}_\theta &= \mathbb{E}[N_\theta](1-c) \quad \forall \theta \in \Theta,\, c = \frac{L_T}{C_1}\left(1 + \max_\theta \frac{\textsc{Conf}_{0,T,\theta}}{\mathbb{E}[N_\theta]}\right) - \max_\theta \frac{\textsc{Conf}_{0,T,\theta}}{\mathbb{E}[N_\theta]}.
\end{aligned}$$

Then, with probability at least $1-\delta$, NL-Guardrail *achieves:*

$$\textsc{Envy}(0) \lesssim L_T \qquad \Delta_{EF}(0, L_T) \lesssim \max\{1/\sqrt{T}, L_T\} \qquad \Delta_{\mathit{efficiency}} \lesssim \min\{\sqrt{T}, 1/L_T\}$$

where $\lesssim$ drops poly-logarithmic factors of T, $o(1)$ terms, and absolute constants.

Proof. The proof follows from Theorem 5.4 by relating $(\Delta_{EF}(\epsilon_1,\epsilon_2), \textsc{Envy}(\epsilon), \Delta_{\mathit{efficiency}})$ to their quantities under the nominal utilities. In particular, we have the following fact, whose proof we defer to the end of this section.

Lemma E.3. *For any allocation X^{alg} we have that:*

$$\begin{aligned}
\textsc{Envy}(0) &\le \textsc{Envy}(\epsilon) + \epsilon\|B\|_\infty\|w'_\theta\|_1 \\
\Delta_{EF}(0,\epsilon) &\le \Delta_{EF}(\epsilon,\epsilon) + \epsilon\|w'_\theta\|_1.
\end{aligned}$$

Using Lemma E.3 we are able to show the claim. Running NL-GUARDRAIL on the *perturbed* utilities, with $\epsilon = L_T$, we have the following upper bounds on $(\textsc{Envy}(L_T), \Delta_{EF}(L_T, L_T), \Delta_{\mathit{efficiency}})$, Theorem 5.4, with probability at least $1-\delta$:

$$\textsc{Envy}(L_T) \leq L_T \qquad \Delta_{EF} \lesssim \max\{1/\sqrt{T}, L_T\}.$$

Plugging this into Lemma E.3, we obtain:

$$\textsc{Envy}(0) \leq \textsc{Envy}(L_T) + L_T\|B\|_\infty\|w'_\theta\|_1 \lesssim L_T + L_T\|B\|_\infty\|w'_\theta\|_1$$
$$\Delta_{EF}(0, L_T) \leq \Delta_{EF}(L_T, L_T) + L_T\|B\|_\infty\|w'_\theta\|_1 \lesssim \max\{1/\sqrt{T}, L_T\}.$$

We conclude the proof of the theorem by noting that the bounds on $\Delta_{\mathit{efficiency}}$ and Δ_{prop} follow trivially from Theorem 5.4 (since these are statements about the allocations themselves, rather than the utilities). □

Proof. We start with $\textsc{Envy}(0)$. Fix all t, t', θ, θ', we have:

$$\begin{aligned}
u(X_{t',\theta'}, \theta) - u(X_{t,\theta}, \theta) &= u_\epsilon(X_{t',\theta'}, \theta) - u_\epsilon(X_{t,\theta}, \theta) - \epsilon\langle w'_{\theta'}, X_{t',\theta'}\rangle + \epsilon\langle w'_\theta, X_{t,\theta}\rangle \\
&\leq u_\epsilon(X_{t',\theta'}, \theta) - u_\epsilon(X_{t,\theta}, \theta) + \epsilon\langle X_{t,\theta}, w'_\theta\rangle \\
&\leq u_\epsilon(X_{t',\theta'}, \theta) - u_\epsilon(X_{t,\theta}, \theta) + \epsilon\|w'_\theta\|_1\|B\|_\infty \\
&\leq \textsc{Envy}(\epsilon) + \epsilon\|w'_\theta\|_1\|B\|_\infty,
\end{aligned}$$

where the final inequality uses the fact that the allocations are upper bounded by B, by the feasibility constraint. The bound on $\textsc{Envy}(0)$ follows by taking the max over t, t', θ, θ'.

We next consider $\Delta_{EF}(0, \epsilon)$. Then, for t, θ, we have that:

$$\begin{aligned}
u(X_{t,\theta}, \theta) - u(X_{t,\theta}^{opt(\epsilon)}, \theta) &= u_\epsilon(X_{t,\theta}, \theta) - u_\epsilon(X_{t,\theta}^{opt(\epsilon)}, \theta) + \epsilon\langle X_{t,\theta}^{opt(\epsilon)} - X_{t,\theta}, w'_\theta\rangle \\
&\leq u_\epsilon(X_{t,\theta}, \theta) - u_\epsilon(X_{t,\theta}^{opt(\epsilon)}, \theta) + \epsilon\|B\|_\infty\|w'_\theta\|_1 \\
&\leq \Delta_{EF}(\epsilon, \epsilon) + \epsilon\|B\|_\infty\|w'_\theta\|_1.
\end{aligned}$$

The result then follows again by taking the max over t and θ. □

F Numerical Experiments: Non-Linear Utilities

In our second set of results we explore the non-linear utility extension to our work (as before, omitting perishing considerations), adopting a similar setup as Sinclair et al. (2022). Namely, we consider the setting with $K = 5$ resources (corresponding to cereal, pasta, prepared meals, rice, and meat) and three types Θ (corresponding to vegetarians, omnivores, and "prepared-food only" individuals). We consider Leontief and Cobb-Douglas utility functions, given by:

- **Leontief**: $u(x, \theta) = \min_{k\in[K]} \frac{x_k}{w_{\theta,k}}$
- **Cobb-Douglas**: $u(x, \theta) = \prod_{k=1}^K w_{\theta,k} x_k^{\alpha_{\theta,k}}$ where $\alpha_{\theta,k} = \frac{1}{K}$

The weights $w_{\theta,k} = p_k \mathbb{1}\{\text{type } \theta \text{ uses resource } k\}$, and p_k are the historical "prices" used by Feeding America in the non-monetary market used to distribute goods to food banks on a national scale (Prendergast, 2017). For example, $w_{\theta,k} = 0$ for the vegetarian type and the meat resource (see Table 2 for the full table of weights).

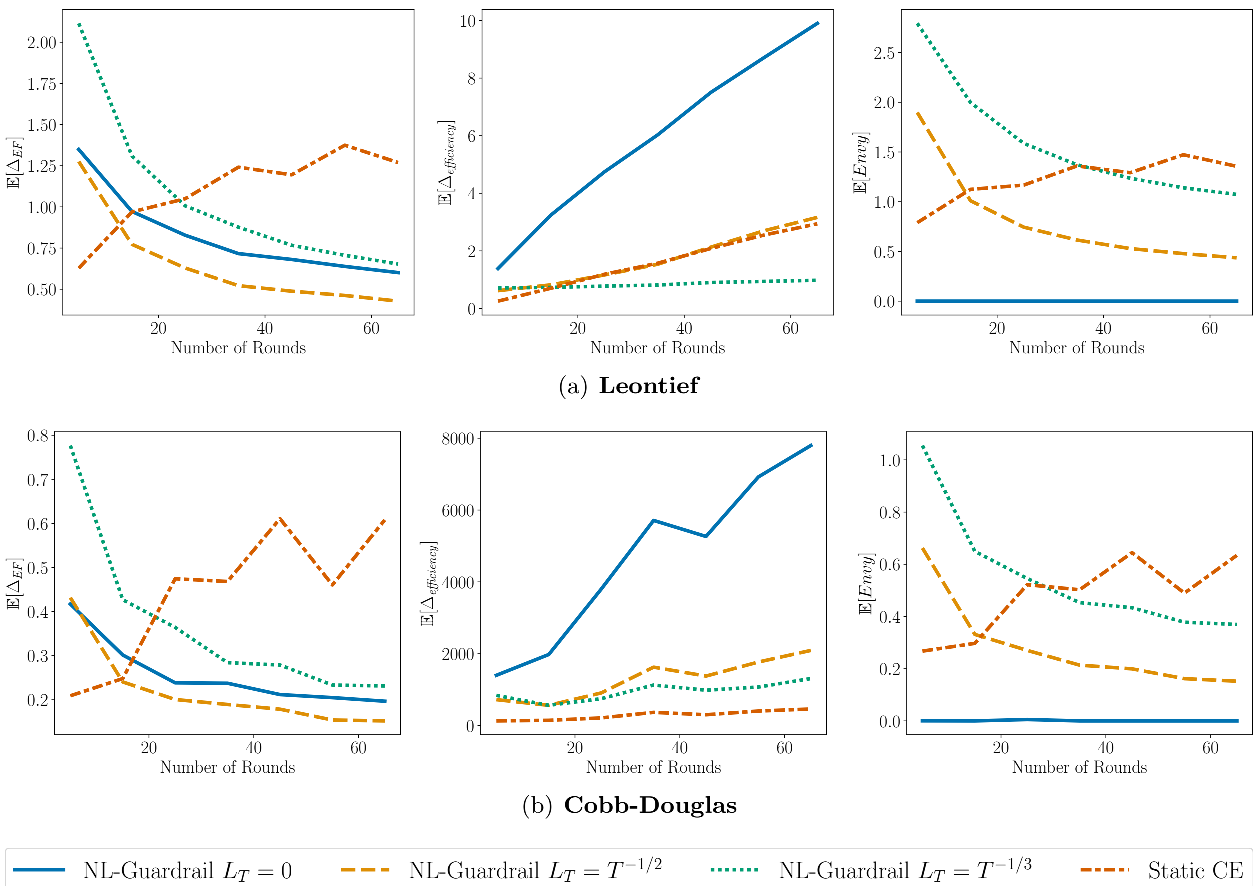

(a) **Leontief**

(b) **Cobb-Douglas**

Figure 2: *Comparison of* NL-Guardrail *for* $L_T \in \{0, T^{-1/2}, T^{-1/3}\}$, *and* Static CE *under Leontief and Cobb-Douglas utilities as we vary* T.

Results.

In this setting we compare the performance of NL-Guardrail on $L_T \in \{0, T^{-1/2}, T^{-1/3}\}$ to Static CE (recall, we've assumed $P_t = 0$ for all t). Our results can be found in Fig. 2.

As in the previous setting, we observe the envy-efficiency trade-off achieved by our algorithm, with better efficiency for larger values of L_T, but worse performance in terms of Δ_{EF}. We again see that the NL-Guardrail algorithms (with $L_T > 0$) outperform NL-Guardrail (with $L_T = 0$) in terms of efficiency, as our algorithms greedily allocate the upper threshold while ensuring budget compliance. Moreover, under Leontief and Cobb-Douglas utilities, NL-Guardrail with parameter $L_T = T^{-1/2}$ achieves lower counterfactual envy than with $L_T = 0$. This serves in contrast to the results from Sinclair et al. (2022) where it was observed that both algorithms achieved similar counterfactual envy. This discrepancy is due to the way sensitivity in the guardrails is propagated under non-linear utilities. In these models we find that the counterfactual envy is incurred in periods receiving the upper guardrail allocation (whereas with additive utilities the difference was the same).

Finally, while Leontief utilities are utility functions of applied interest, they fail to satisfy the assumptions upon which our theoretical results rely. Our results highlight that the algorithmic guardrail technique is robust to such phenomena. See Appendix E for more discussion on this fact.

Table 2: Weights w_k for the different products. Here we use the weights taken from the historical prices used in the market mechanism to distribute food resources to food pantries across the United States (Prendergast, 2017).

Resource	Cereal	Pasta	Prepared Meals	Rice	Meat
Weights (type $\theta =$ omnivore)	3.9	3	2.8	2.7	1.9
Weights (type $\theta =$vegetarian)	3.9	3	0	2.7	0
Weights (type $\theta =$prepared-only)	3.9	3	2.8	2.7	0

G Simulation Details

We further observe the relationship between these metrics in the radar plots included in Fig. 3.

Experiment Setup: Each experiment was run with 200 iterations where the relevant plots are taking the mean of the related quantities. In all experiments the budget $B = \sum_{t,\theta} \mathbb{E}[N_{t,\theta}]$ so that β_{avg} scales as a constant as we vary the number of rounds T. All randomness is dictated by a seed set at the start of each simulation for verifying results. In the algorithm description we ignore all logarithmic terms, setting their value equal to one. Moreover, since all of the examples for perishing resources satisfy Condition (2) of Theorem 4.1, it is without loss of generality to set $\eta_t = \sqrt{T - t}$ for all $t \in [T]$.

Computing Infrastructure: The experiments were conducted on a personal computer with an AMD Ryzen 5 3600 6-Core 3.60 GHz processor and 16.0GB of RAM. No GPUs were harmed in these experiments.

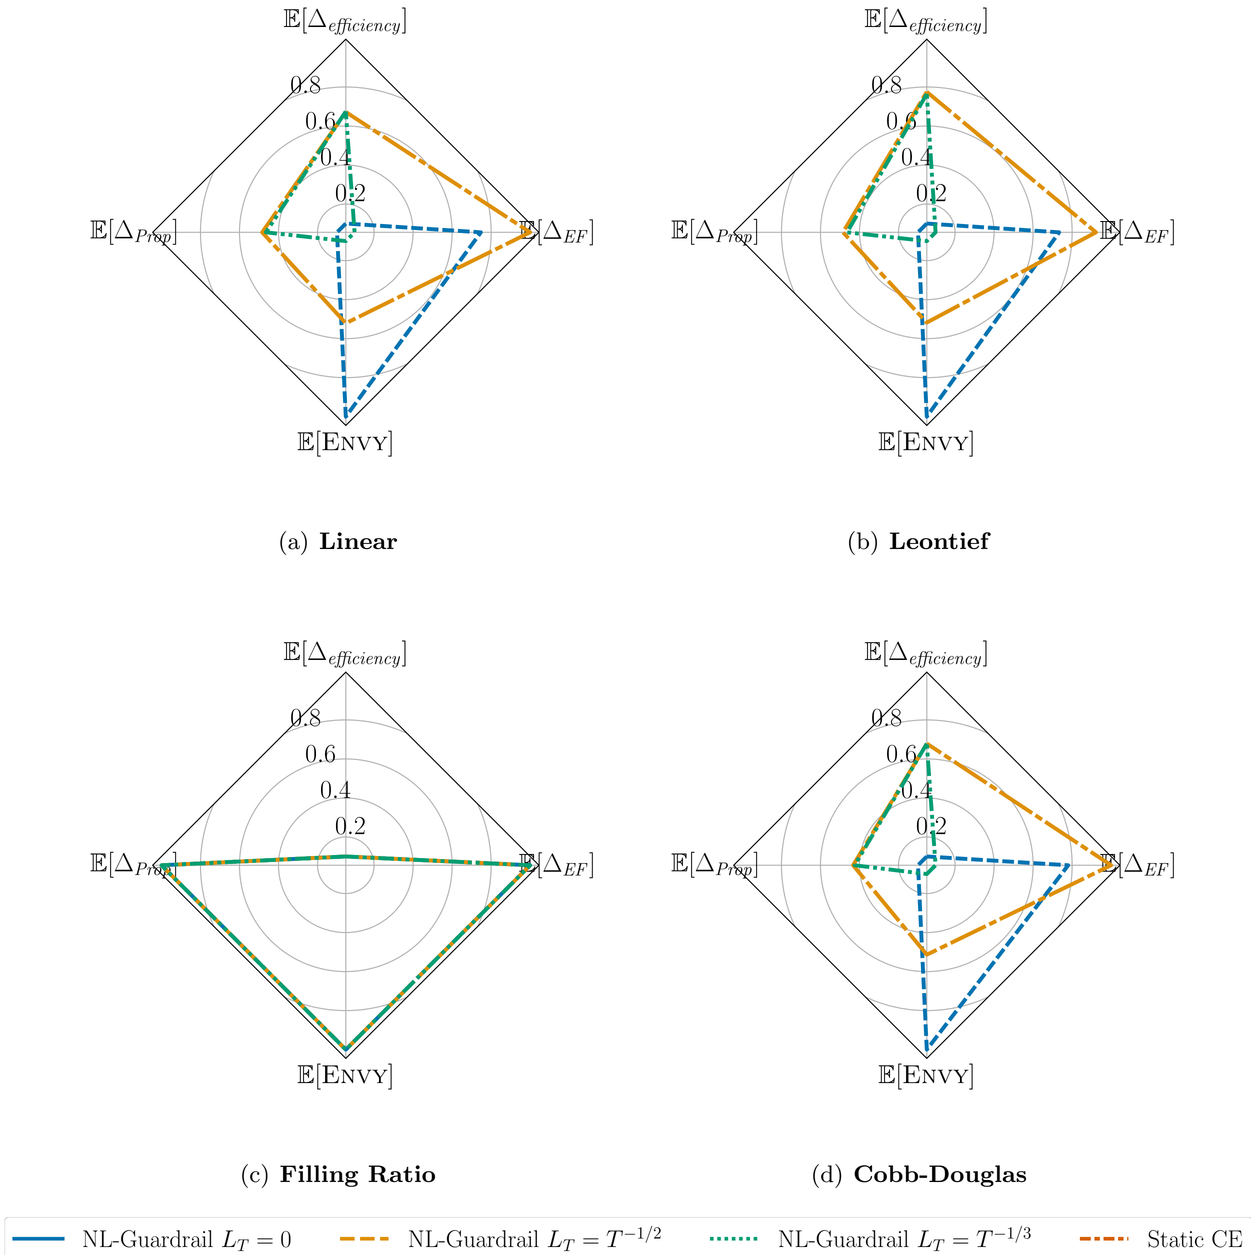

(a) **Linear**

(b) **Leontief**

(c) **Filling Ratio**

(d) **Cobb-Douglas**

Figure 3: Comparison of NL-Guardrail for $L_T = T^{1/2}$ and $L_T = T^{1/3}$ Fixed Threshold, Hope-Online, and Hope-Full under Leontief and filling ratio utilities. The values here are normalized to be between $[0, 1]$ to better highlight the performance of the algorithms, where larger values correspond to better performance.

H Useful Lemmas

We use the following standard theorems throughout the proof. See, e..g. Vershynin (2018) for proofs and further discussion.

Theorem H.1 (Hoeffding's Inequality)**.** *Let* $X_1, \ldots, X_n$ *be independent random variables such that* $a_i \leq X_i \leq b_i$ *almost surely and set* $S_n = \sum_i X_i$*. Then we have that for all* $t > 0$*:*

$$\mathbb{P}(|S_n - \mathbb{E}[S_n]| \geq t) \leq 2 \exp\left(-\frac{2t^2}{\sum_i (b_i - a_i)^2}\right).$$

Theorem H.2 (Chernoff Bound for Sum of Independent Bernoulli Random Variables)**.** *Consider a sequence of Bernoulli random variables* $(X_i)_{i \in [N]}$*, independently distributed with probability of success* $p_i \in (0, 1)$*. Let* $X = \sum_i X_i$*, and let* $\mu = \mathbb{E}[X] = \sum_i p_i$*. Then for all* $0 < \delta < 1$*:*

$$\mathbb{P}(|X - \mu| \geq \delta\mu) \leq 2 \exp\left(-\frac{\mu\delta^2}{3}\right).$$